Task Routing for Prediction Tasks

Haoqi Zhang*, **Eric Horvitz**†, **Yiling Chen***, and **David C. Parkes***
*Harvard SEAS
Cambridge, MA 02138, USA
{hq, yiling, parkes}@eecs.harvard.edu
†Microsoft Research
Redmond, WA 98052, USA
horvitz@microsoft.com

ABSTRACT

We describe methods for routing a prediction task on a network where each participant can contribute information and route the task onwards. *Routing scoring rules* bring truthful contribution of information about the task and optimal routing of the task into a Perfect Bayesian Equilibrium under common knowledge about the competencies of agents. Relaxing the common knowledge assumption, we address the challenge of routing in situations where each agent's knowledge about other agents is limited to a local neighborhood. A family of *local routing rules* isolate in equilibrium routing decisions that depend only on this local knowledge, and are the only routing scoring rules with this property. Simulation results show that local routing rules can promote effective task routing.

Categories and Subject Descriptors

J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms

Algorithms, Economics, Theory

Keywords

Scoring rules, task routing, social networks

1. INTRODUCTION

Organizations rely on a mix of expertise and on means for identifying and harnessing expertise for completing different kinds of tasks. The ability to leverage the expertise and interests of individuals effectively is crucial for the success of an organization. Accomplishing a task may require the expertise of multiple actors, and harnessing that expertise requires identifying who the experts are and providing proper incentives for inducing contributions.

One approach to coordinating expertise is to pool knowledge about competencies and preferences and to assign tasks in a centralized manner. Another approach is to rely on individuals distributed across an organization to select tasks themselves. Both approaches have flaws. In the former, an organization or system may not know which individuals have the required expertise. In the latter, while individuals may often be able to gauge their own expertise, they may not know which tasks best match their respective competencies.

In social networks and organizations, an individual's knowledge extends beyond their own expertise on tasks and topics to knowledge about the expertise of others. For example, members within the same research group know whom within that group can best review a paper, or best contribute to answering a research question. Members of a social network may know who among their friends can best answer a particular question, or otherwise provide valuable opinions on a topic of discussion. Even in situations where an individual cannot identify an expert who can best contribute to a task, they may know others who would likely know experts, or be able to identify subsets of individuals among whom the requisite expertise is likely to exist (e.g., people who share a particular interest).

We explore principles and methods for *task routing* that aim to harness the ability of people or automated agents to both contribute to a solution, and to route tasks to others who they believe can also effectively solve and route. Task routing provides an interesting paradigm for problem solving in which individuals become engaged with tasks based on their peers' assessments of their expertise. On the task level, effective task routing aims to take advantage of agents' knowledge about solving problems as well as agents' knowledge about other agents' abilities to contribute. Agents make routing decisions in a peer-to-peer manner, and the system rewards participating agents for their contributions. On the organizational level, task routing may provide a means for bringing tasks to individuals effectively, where agents' routing decisions take into account not only an individual's expertise on the particular task, but also their ability to contribute as a router.

Methods for automated and manual routing of tasks have been employed in online networks. For example, question-answering services such as Aardvark [10] allow a user to ask questions in natural language, which the system interprets and automatically routes to appropriate individuals in the user's social graph based on an assessment of who is best able and willing to provide an answer. Aardvark also allows for peer routing, where a user can manually route questions to others, enabling the system to reach users outside its fund of knowledge about people and their expertise.

We consider methods for routing and solving tasks with a focus on the challenge of efficiently obtaining accurate probability assessments about an uncertain event. For this task,

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

a question is passed among individuals on a network, and each participant can update the posterior probability and forward the task to a neighbor. We introduce *routing scoring rules* for incentivizing contributions. Given an assumption of common knowledge about the amount of information held by each agent on the network, we obtain truthful reporting of posterior probability assessments and optimal routing in a Perfect Bayesian Equilibrium. Even with this common knowledge assumption, we find that the equilibrium strategy on a general network requires finding a routing path through individuals that have the most information in aggregate and is NP-hard. In comparison, a myopic routing rule is optimal on simple topologies such as cliques.

A second difficulty is that common knowledge is unlikely to hold for large social networks where each agent's information about the competencies of others is limited to a local neighborhood (e.g., friends, and perhaps friends of friends). To handle such cases, we also consider task routing where knowledge about others' abilities may be limited to only those agents in an agent's local neighborhood. Unfortunately, equilibrium routing under the routing rules becomes even more computationally challenging, requiring an agent on the path to perform inference that takes into consideration the previous routing decisions of agents.

Beyond formalizing the joint routing and solution challenge, our main contribution is to introduce a family of *local routing rules* that isolate simple routing decisions in equilibrium, while still taking advantage of knowledge about the expertise of others to promote effective routing decisions. We achieve this by incentivizing agents to make routing decisions based on short, locally optimal paths that can be computed easily using shared local knowledge. In summary, *we design incentive schemes that explicitly enable equilibrium behavior for which the inference required of agents is tractable.*[1] We provide a full characterization of local routing rules, and show that they are the only routing scoring rules that induce truthful equilibria in which agents' routing decisions can be computed using only local common knowledge. Simulation results demonstrate that equilibrium routing strategies based on local routing rules lead to effective information aggregation.

1.1 Related work

Leveraging individuals' abilities to both solve and route is a key component of the winning team's strategy in the DARPA Red Balloon Challenge [13]. The task was to find large, salient helium-filled balloons placed in ten undisclosed locations across the continental United States. The winning team introduced an incentive mechanism that uses a limited budget to incentivize individuals to look for balloons and to let their friends know about the task; see also Emek et al. [8] and Douceur and Moscibroda [7] for related theoretical analysis, and work on query incentive networks [12, 2, 5] that analyze games in which players split rewards to recruit others to answer a query. The mechanism used by the winning team differs from those in our work because it aims to induce agents to broadcast the task to everyone they know regardless of their expertise or knowledge of others' expertise, whereas mechanisms in our work aim to induce agents to identify particular experts that can best contribute to the task and route it to others.

[1] This is analogous to the role of strategyproofness in simplifying strategic problems facing agents in mechanism design.

The problem of task routing is also related to the problem of decentralized search on networks in which the goal is to find a target node quickly through local routing decisions [15, 6, 16, 11, 1]. In such work, the goal is to identify a *single* target node representing a particular individual; while a single-target task differs from the task routing problem we seek to solve, the results on its solution provides theoretical and experimental support for the prospect that routing decisions with local information may have effective, global performance.

One can view routing scoring rules as an extension of market scoring rules [9] used in prediction markets. Market scoring rules provide proper incentives for individuals to improve probability estimates by contributing additional information. The major difference between task routing and a prediction market is in the 'burden' of identifying expertise: while prediction markets place the responsibility on individuals to find prediction tasks for which they have useful information, task routing incentivizes individuals to notify others with appropriate expertise who may otherwise be unaware of the task.

2. MODEL

To formalize the setting, consider a single prediction task T, for which we would like to gather an accurate probability assessment of the true state $\omega \in \Omega$. The probability assessment task can be for any state of the world that will be revealed later in time, e.g., "Will it snow next Tuesday in Boston?" or "Will the Celtics win the NBA championship this year?" We consider discrete state spaces, and assume without loss of generality a binary state space, such that $\Omega = \{Y, N\}$.

Consider a game with n players, where each player is represented by a node on the *routing graph* $G = (V, E)$. Edges in the graph may be directed or undirected, and indicate whether a particular player can route the task to another player. The task is initially assigned to a source player named player 1, with later players on a routing path numbered sequentially. The source player is either determined by the system, or the individual who originally posed the task. The source player is asked to update the probability of state Y from the prior probability p^0 to some probability p^1, and, in addition, to route the task to a neighbor. The selected neighbor is then asked to update the assessment p^1 to p^2 and route the task to a neighbor, and so on, until the game ends after a pre-specified number of rounds R that denotes when a final assessment must be made. We assume players receiving the task are provided with a list of people or agents who have participated so far, as well as information about the number of rounds that remain. *Our goal is to design incentive mechanisms that will induce each player to update probability assessments truthfully and route the task to other players that can best refine the prediction, so as to arrive at an accurate assessment after R rounds.*

We model players' knowledge about the task as follows: the true state of the world is drawn according to the probability distribution $\Pr(Y) = p^0$ and $\Pr(N) = 1 - p^0$, which is common knowledge to all players. While no player observes the true state directly, each player may receive additional information about the true state. To model this state of affairs, each player privately observes the outcome of some number of coin flips drawn according to a commonly known distribution that depends on the true state. Different players may observe different numbers of coin flips, where players observing more coin flips are *a priori* more knowledgeable.

Formally, we represent player i's signal c_i as a random bit vector of length l_i, where bit c_{ik} is a random variable over the outcome of the k-th coin flip observed by player i. We assume the value of bits of signal are *conditionally independent* given the true state, and drawn from the same distribution (known to all players) for all players and all bits, such that $\Pr(c_{ik} = H|\omega) = \Pr(c_{jm} = H|\omega)$ for all players i, j, bits k, m, and realization H (head). Each bit of signal is assumed to be *informative*, that is, $\Pr(c_{ik} = H|\omega = Y) \neq \Pr(c_{ik} = H|\omega = N)$ for all i, k. We also assume that bits of signal are *distinct*, that is, $\Pr(\omega = o|c_{ik} = H) \neq \Pr(\omega = o|c_{ik} = T)$ for all i, k, o, where H is heads and T is tails.[2] We assume the realization of each player's signal is private, but make different assumptions about the knowledge of one player about the *number* of coin flips of another player.

With conditionally independent signals, each player can properly update the posterior probability without having to know the signals of previous players or their length, as long as previous updates were done truthfully [4]. This is useful practically in that players do not have to keep track of nor communicate their signals, and can simply report an updated posterior probability that sufficiently summarizes all information collected thus far.

3. ROUTING SCORING RULES

With rational, self-interested players who have no intrinsic value (or cost) for solving or routing a particular task, effective task routing requires mechanisms that will incentivize players to both truthfully update and report posterior probabilities and to route tasks to individuals who can best refine the predictions of the tasks. In this section, we review strictly proper scoring rules and market scoring rules for incentivizing truthful reports, and introduce routing scoring rules, which also incentivize effective routing decisions.

In the forecasting literature, *strictly proper scoring rules* [14] are mechanisms that strictly incentivize a forecaster to truthfully reveal his subjective probability of an event, typically under the assumption that agents are risk neutral. The outcome of the event is assumed observable in the future, and payments are conditioned on the outcome. A well-known strictly proper scoring rule is the *quadratic scoring rule*, under which a player reporting probability q for state Y is rewarded $1 - (1 - q)^2$ when the true state is Y and $1 - q^2$ when the true state is N. Other strictly proper scoring rules include the logarithmic and spherical scoring rules, and any strictly proper scoring rule can be scaled or normalized via linear transformations to form another strictly proper scoring rule [3].

Market scoring rules [9] extend strictly proper scoring rules to settings where we wish to aggregate information across multiple people. Given a sequence of reports, player i reporting p^i is rewarded $s_i - s_{i-1}$, where s_i denotes the score of player i as computed by some strictly proper scoring rule applied to this agent's report alone. Note that since strictly proper scoring rules incentivize accurate reports, a player's reward under a market scoring rule is positive if and only if he improves the prediction.

[2]These assumptions rule out degenerate cases and can be made without loss of generality. A signal that is not informative can be removed from the signal space, and two signals that are not distinct can be treated as the same signal.

Building on market scoring rules, we introduce *routing scoring rules* to incentivize accurate predictions, along with effective routing decisions.

DEFINITION 1. *A* **routing scoring rule** *defines a sequence of positive integers $k_1, \ldots, k_{R-1}$, which rewards players $i \in \{1, \ldots, R-1\}$ on the routing path:*

$$(1-\alpha)s_i + \alpha s_{i+k_i} - s_{i-1}$$

where s_i is the score under an arbitrary strictly proper scoring rule, $\alpha \in (0, 1)$ is a constant, and $i + k_i \leq R$ for all players i. Player R reports but does not route and is paid $s_R - s_{R-1}$.

In a routing scoring rule, player i's payment is based on the marginal value the player provides for refining the prediction, as measured by his report and the report of the player who receives the task k_i steps after him, in comparison to the report of the player just before him. For player 1, s_0 denotes the score computed with respect to the prior p^0. Each player i can be paid for up to $R - i$ steps forward, and the final player R does not route and is paid by the market scoring rule $s_R - s_{R-1}$.

Intuitively, routing scoring rules reward players who are experts and players who are knowledgeable about the expertise of other players. We introduce here several routing scoring rules of particular interest. We first consider the *myopic routing scoring rule* (MRSR), which sets $k_i = 1$ for all players $i < R$. This routing scoring rule aims to reward a player for submitting accurate probability assessments and routing in a greedy manner to the player who can most accurately refine the probability assessment.

LEMMA 1. *The total payment from the system in the routing game with MRSR is $s_R - s_0 + \alpha(s_R - s_1)$.*

The lemma follows from taking telescoping sums, and states that, for MRSR, the center needs to only pay for the difference between the final assessment and the initial assessment, since each player is only paid for the additional information they provide and their routing decision.

We can extend the MRSR to reward players' routing decisions based on the accuracy of information after $k_i = \min(k, R - i)$ more players have provided their information. The *k-step routing scoring rule* (kRSR) rewards a player based on his report, as well as the eventual consequence of his routing decision k steps into the future. Unlike MRSR, kRSR rewards players for routing to players who may not have information themselves but who are still able to route to others who do.

In particular, when player i's routing payment is based on player R's score, that is, $i + k_i = R$, for all i, we call this the *path-rewarding routing scoring rule* (PRSR). As its name suggests, this routing scoring rule seeks to focus a player's attention on the final consequence of his routing decision, as judged at the end of the solving and routing process.

The choice of routing scoring rule affects players' routing decisions in equilibrium, which in turn affects how much information is aggregated. To see the connection between a player's score and the amount of information aggregated, note that the expected score is strictly increasing in the total number of coin flips collected:

LEMMA 2. *Let S' and S'' denote two possible sequences of players through the first k rounds of the routing process that* *are identical up to player $i < k$. Assume all players truthfully update posterior probabilities, and that player i knows l_j for players $i < j \leq k$ on S' and S''. Let $E^i_S[s_k]$ denote player i's ex-ante expectation of the score after player k's report in path S. $E^i_{S'}[s_k] > E^i_{S''}[s_k]$ holds if and only if $\sum_{m \in u(S')} l_m > \sum_{n \in u(S'')} l_n$, where $u(S)$ is the (unique) set of players in S.*

PROOF. (sketch) Assume without loss of generality that there are a total of n coin flips in S', and $n + m$ coin flips in S'', $m > 0$. The expected score of player k from S'' consists of two (hypothetical) parts, (a) the score he would get when giving a prediction after receiving the first n coin flips, denoted $s_{[n]}$, and (b) the difference in the score he would get by changing his prediction after receiving the next m coin flips, denoted $s_{[n+m]} - s_{[n]}$. The expectation of the first part is the same as the expected score of player k from S', and the expectation of the second part is always non-negative given any strictly proper scoring rule. □

Intuitively speaking, additional bits of information can only improve the accuracy of the prediction in expectation. Since strictly proper scoring rules reward accuracy, collecting more coin flips will lead to higher scores in expectation.

4. CASE OF COMMON KNOWLEDGE

Having introduced routing scoring rules of interest, we consider an equilibrium analysis of the associated routing game. We first consider the case where the number of coin flips l_i observed by each player i is common knowledge.[3] Note the actual signal realizations are still assumed private.

4.1 Clique topology

Let us now consider the routing game on a *clique*, where each player can route the task to any other player. From Lemma 2, and given the clique topology, an optimal routing algorithm can just route myopically and collect as many coin flips as possible at each step. This is because there is no opportunity cost for being greedy in this way, due to the clique topology. We have the following equilibrium result:

THEOREM 1. *Assume the number of coin flips of each player is common knowledge, and that players are risk neutral. Consider a routing game in which the routing graph is a clique, and let $S_{>i}$ denote the set of players who have yet to receive the task after i rounds. Under the myopic routing scoring rule, it is a Perfect Bayesian Equilibrium (PBE) for each player i to truthfully update the posterior probability, and to route the task to player $i+1 \in \text{argmax}_{m \in S_{>i}} l_m$, with the belief that all other players update the posterior probability truthfully.*

PROOF. (sketch) We show that no player wishes to deviate from the equilibrium strategy, given the belief that all other players report truthfully. Consider player i. To prove the theorem, we first show that player i should honestly update the posterior beliefs by establishing that (a) truthful reporting maximizes s_i, and that (b) for any player m who may be routed the task, truthful reporting by player i maximizes the score s_m. For (a), note that, since s_i is based on a strictly proper scoring rule, truthful reporting maximizes the expectation of s_i. For (b), note that the expected score of s_m (from the perspective of player i) is strictly greater when player i reports honestly because s_m is based on a strictly proper scoring rule. It is left to show that player i maximizes s_{i+1} by routing to the player in $S_{>i}$ with the most coin flips; this follows from Lemma 2. □

4.2 General networks

We now turn to consider routing games on general networks, with missing edges; e.g., only managers can route tasks between teams, only professors can route questions to other professors, and only friends can route to friends.

We can state the algorithmic problem of finding the optimal route in terms of collecting coin flips:

PROBLEM 1. *Consider the routing graph $G = (V, E)$, where nodes are assigned non-negative integer weights w_i (coin flips). Given a starting node o, find a path of length at most k such that the sum of weights on the path is maximized.*

Note that a player can route to another player who have received the task before (e.g., the path need not be *simple*), but no additional information is collected in subsequent visits.

Immediately, we see that myopic routing will not always find the optimal solution to this problem, as routing to the neighbor with the most coin flips does not consider the consequence on future routing decisions and can now convey an opportunity cost.

We can show that this problem is NP-hard for variable path length k:

LEMMA 3. *Problem 1 is NP-hard.*

PROOF. Consider a reduction from the Hamiltonian Path problem. Let all nodes have weight 1, and set $k = |V|$. The solution path has total weight $|V|$ if and only if all nodes are visited within k steps, that is, a Hamiltonian Path exists. □

While the problem is NP-hard for a variable path length k, for small constant k the optimal path may be tractable to compute via exhaustive search.

Intractability is not the only difficulty faced. Even if players can compute the optimal path, we still need to find incentives that induce players to honestly report their information and to route along the optimal path. The path-rewarding routing scoring rule does just that.

THEOREM 2. *Assume the number of coin flips of each player is common knowledge, and that players are risk neutral. Let $S_{>i}$ denote the set of players who have yet to receive the task after i rounds. Let Q_i denote a solution to problem 1 for which $k = R - i$, $o = i$, and $w_m = l_m$ if $m \in S_{>i}$ and 0 otherwise. Under the path-rewarding routing scoring rule, it is a PBE for each player i to truthfully update the posterior probability, and to route the task to the next player in the path provided by Q_i, with the belief that all other players follow this strategy.*

Since PRSR rewards each agent's routing decision based on the final score, it is in each agent's interest to maximize the number of coin flips collected along the entire routing path. We can show that reporting honestly and routing this way is the only behavior that can be supported in equilibrium under PRSR:

[3]By taking appropriate expectations, the analysis throughout the paper extends easily to settings where players are equally well-informed but are uncertain about the number of coin flips that other players observe.

Theorem 3. *The set of PBE identified in Theorem 2 (corresponding to possible ties in the solution to problem 1) are the only PBE of the routing game under PRSR.*

Proof. (sketch) Given any routing path, by backward induction every player should update the posterior probability truthfully because agents' scores are computed using a strictly proper scoring rule. Given that players update truthfully, by backwards induction every player i should route along the path identified by some solution Q_i because maximizing the number of coin flips collected maximizes the routing portion of each player's score (Lemma 2). □

5. LOCAL COMMON KNOWLEDGE

Although people may know one another's expertise in small organizations, the common knowledge assumption becomes unreasonable for larger organizations and social networks. Any given individual will not necessarily know everyone else, and may only have summary information about the expertise and connectivity of individuals outside of a local neighborhood.

We replace the common knowledge assumption with a requirement that individuals all attain the same minimal level of knowledge about each others' expertise within a particular size of local neighborhood, defined by the number of hops between agents. For example, all friends of a particular person are aware of his expertise, and friends of his friends may also be aware; people may know a local portion of the routing graph, e.g., individuals typically know not only their friends but also their friends' friends.

Definition 2. *A routing game satisfies the* **local common knowledge assumption within m-hops** *if, for all nodes (individuals) i, (a) l_i is common knowledge to all individuals connected to i via some path of length at most m, and (b) i knows all paths of length at most m connecting i to other individuals, and this is common knowledge.*

For example, 1-hop local common knowledge assumes all friends of a particular person know the person's level of expertise, and 2-hop local common knowledge extends this shared knowledge to his friends of friends. Note that the local common knowledge assumption within m-hops is just a minimal requirement, and does not preclude a player having more information.

Given that a player may only have m-hop local common knowledge, let's consider the problem facing such a player in deciding how to route to maximize the final prediction quality after R steps. Routing optimally may require a player to use the history of routing decisions to infer why certain people were not routed the task (but could have been), based on which to perform inference about the amount of information of different agents in the network. Furthermore, optimal routing requires a player to make inferences about the value that can be generated from the routing decisions of subsequent players beyond his locality. Not only is such reasoning complex and likely impractical, any equilibrium to induce optimal routing is likely to be fragile as it requires players to adopt priors on other players' beliefs.

An attempt to avoid such issues may suggest incentivizing players based on a m-step routing rule whenever the local common knowledge assumption holds for m-hops. The problem with this suggestion is that a player still has to consider routing decisions of players outside its locality because maximizing its payoff requires considering the routing decisions of the chain of players within its locality. For example, consider the two-step routing rule (see bottom of Figure 1). For any player, the score two steps forward will depend in part on the routing decision of the next player. But since the next player is also paid by the two-step routing rule, his routing decision will depend not only on the amount of information held by the player after him, but also that player's routing decision. Since each player has to consider the routing decision of the next player, each player has to reason about the future routing decisions of all players down the routing path, just to compute the expected score after two steps.

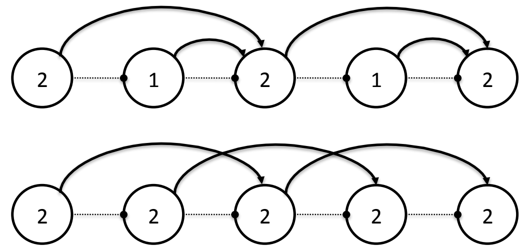

Figure 1: Illustration of the 2-1-2-1 and 2-step routing rules. Arrows depict dependencies in routing payments.

This motivates the family of *local routing rules*, under which players' strategies in equilibrium rely only on computations based on local information, but nevertheless take advantage of the available local common knowledge. We define the notion of a local strategy as follows:

Definition 3. *A player i in a routing game adopts a* **m-local strategy** *if its routing decision depends only on m-hop local common knowledge and is invariant to any beliefs the player might have about players outside of its own locality.*

Let us first consider the following local routing rule, designed to be useful with 2-hop local common knowledge:

Definition 4. *The* **2-1-2-1 routing rule** *is a routing scoring rule which sets $k_i = 2$ if i is odd and $i < R-1$, and $k_i = 1$ otherwise.*

The 2-1-2-1 routing rule incentivizes players to compute locally optimal paths of length two (see top of Figure 1), which can be computed with local common knowledge, and so inference is not required. As even players are paid based on the myopic routing scoring rule, they will route to the available player with the most number of coin flips. Since each odd player knows the number of coin flips that can be collected from the next even player and also from the odd player that is then routed the task, he can compute the best local path without regard to routing decisions beyond his locality. Note that players still need to take into account which other players have already participated, but no other inference based on history is necessary.

Expanding on the idea, we construct a class of routing scoring rules (e.g., MRSR, 2-1-2-1, 3-2-1-3-2-1, ...) that incentivize players to compute *locally optimal paths* for m-hop local common knowledge:

Definition 5. *The* **m-hop routing rule** *is a routing scoring rule which sets $k_i = \min[m - (i-1) \bmod m, R - i]$.*

We can characterize the equilibrium behavior as follows:

THEOREM 4. *Assume that players are risk neutral and m-hop local common knowledge holds. Let $S_{>i}$ denote the set of players who have yet to receive the task after i rounds. Let Q_i denote a solution to problem 1 for which $k = \min[m-(i-1) \bmod m, R-i]$, $o = i$, and $w_j = l_j$ if $j \in S_{>i}$ and 0 otherwise. Under the m-hop routing rule, it is a PBE for each player i to truthfully update the posterior probability, and to route the task to the next player in the path provided by Q_i, with the belief that all other players follow this strategy.*

PROOF. (sketch) Using similar arguments as the proof sketch for Theorem 1, we can show that players should truthfully update the posterior probability. To show player i should route based on Q_i, we first note that Q_i is computable given m-hop local common knowledge. Since Q_i maximizes the number of coin flips collected in the next k steps, Lemma 2 proves the point, and the theorem. □

The main idea behind the m-hop routing rule is that each player can compute his best routing action with respect to the decisions in his locality and without regard to routing decisions beyond his locality. This property can be satisfied by other local routing rules as well. For example, when $m = 3$, the 3-1-1-3-1-1 routing rule is one in which the first of three players in sequence is paid by the score three steps forward, but in which the next two players are each paid myopically. Note that here the first player can still compute its optimal routing decision using only local common knowledge, by computing the routing decisions of others in its locality via backwards induction. We can thus characterize the entire family of local routing rules:

DEFINITION 6. *Given m-hop local common knowledge, the family of* **local routing rules** *contains routing scoring rules $k_1, \ldots, k_{R-1}$ that satisfies* **local reasoning***, that is, $k_{i+j} + j \leq m$ for all i and $0 \leq j < k_i$.*

The local reasoning condition ensures that local routing rules can only reward players whose routing decisions may affect the payoff of an earlier player based on the routing decisions of future players that are within m hops of that earlier player. In other words, it considers the set of routing scoring rules for which the payment to any player should only depend on the local information that player is guaranteed to hold. For example, the 2-1-2-1 routing rule satisfies local reasoning for $m = 2$ because for an odd i, $k_i \leq 2 \leq m$ and $k_{i+1} + 1 = 2 \leq m$, and for an even i, $k_i = 1 \leq m$. However, the two-step routing scoring rule violates local reasoning, because for all $i < R - 2$, $k_{i+1} + 1 = 3 > m$. Note that the m-hop routing rule satisfies local reasoning, since k_i is set such that $k_{i+j} + j = m$ for all appropriate i and j in the above definition.

We argue that using a local routing rule is necessary and sufficient for the existence of an equilibrium in which agents follow m-local, truthful strategies. We first show sufficiency:

THEOREM 5. *Assume that risk neutrality and m-hop local common knowledge holds. For any node i and possible path $n_{i+1}, \ldots, n_{i+k_i}$ from i, let the weights w_j on node j be l_j if j has yet to be visited up until then, and 0 otherwise. For any local routing rule, consider the following dynamic program:*

$$V(n_{j+1}, \ldots, n_{j+k_j} | n_1, \ldots, n_j) = \max_{j+1,\ldots,j+k_{j+1}} [\sum_{b=1}^{k_{j+1}} w_{j+b} + V(n_{j+k_{j+1}+1}, \ldots, n_{j+k_j} | n_1, \ldots, n_{j+k_{j+1}})]$$

$$V(\emptyset | n_1, \ldots, n_{j+k_j}) = 0 \quad \forall n_1, \ldots, n_{j+k_j}$$

Let $n^*_{i+1}, \ldots, n^*_{i+k_i} = \operatorname{argmax} V(n_{i+1}, \ldots, n_{i+k_i} | n_1, \ldots, n_i)$ denote a solution of the dynamic program. It is a PBE for each player i to truthfully update posterior probabilities and to route the task to n^*_{i+1}, with the belief that all other agents follow this strategy.

PROOF. (sketch) To prove the theorem, we first note that all players would truthfully update the posterior probability along the path as we had previously argued, as doing so maximizes the scores computed, based on its assessment and based on the assessments collected from those routed the task via the routing payment. Second, as the variables and parameters of the dynamic program are only the nodes in paths of length at most k_i from i, and by the local reasoning assumption $k_i \leq m$, players follow m-local strategies in which the information that each player i needs to compute the dynamic program is within m hops and thus known to player i. Finally, given the routing decisions of others down the path, the number of coin flips collected is by definition maximized by the routing decisions along the computed path. Applying Lemma 2 proves the point, and the theorem. □

THEOREM 6. *For any local routing rule, the set of PBE identified in Theorem 5 (corresponding to possible ties in the optimal solution to the dynamic program) are the only PBE of the routing game under that local routing rule.*

THEOREM 7. *The only routing scoring rules that induce for every routing game a truthful PBE (where players honestly update probability assessments) in m-local strategies are local routing rules.*

PROOF. (sketch) Assume for sake of contradiction that there exists a routing scoring rule that induces a truthful PBE for all routing games in m-local strategies but is not a local routing rule. Since this routing scoring rule violates local reasoning, there must be some i in the sequence for which there exists some j such that $k_{i+j} + j > m$, $0 \leq j < k_i$. Consider the first such i and j.

First consider the case where $j = 0$. We construct a graph with two paths (top and bottom), as is shown in Figure 2:

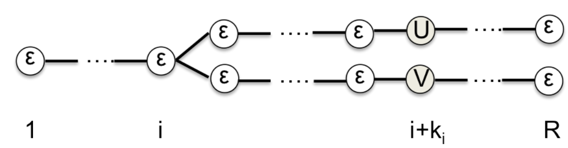

Figure 2: Routing game construction for $j = 0$ case.

Based on the construction, consider two routing games G and G', where in game G the coin flips held by U and V are 1.5ϵ and 1.6ϵ respectively and in game G' the coin flips at U and V are reversed. Due to the violation of local reasoning at i for $j = 0$, by construction U and V are more than m hops from player i. In a PBE with m-local strategies, it is thus necessary for the routing decisions of player i to be independent of the number of coin flips held by players at U and V, that is, for the routing decision to be the same for these two games G and G'.

We show that player i's best response to the equilibrium strategies of the other agents depends on G or G'. For both games, using backwards induction, all players strictly prefer to route the task forward (to the right) instead of backwards at any given point in time and for any lookahead depth as induced by their routing payment, because its expected payment is based on the number of coin flips collected and one can always collect more coin flips in the forward direction (because for any player, going backwards would necessitate visiting a node that's been visited before and thus has no new coin flips to share). Since in game G player i would collect more coin flips by routing up due to the higher value at U over V and the reverse is true in game G', player i's best response would be different, which contradicts our assumption.

Now consider the case where $j > 0$. We construct a graph with three paths (top, middle, and bottom), as is shown in Figure 3:

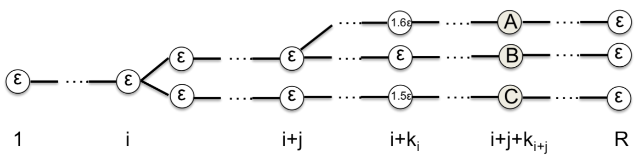

Figure 3: Routing game construction for $j > 0$ case.

Based on the construction, consider two routing games G'' and G''', where in game G'' the coin flips held by A, B, and C are ϵ, ϵ, and ϵ respectively, and in game G''' are ϵ, 1.7ϵ, and 1.7ϵ, respectively. Due to the violation of local reasoning, by construction A, B, and C are more than m hops from player i. In a PBE with m-local strategies, it is thus necessary for the routing decisions of player i to be independent of the number of coin flips held by players at A, B, and C, that is, for the routing decision to be the same for G'' and G'''.

We show that player i's best response to the equilibrium strategies of the other agents depends on G'' or G'''. We first consider game G''. Using backwards induction, note that each player must strictly prefer to route the task forward (to the right) instead of backwards at any given point in time, regardless of the lookahead induced by their routing payment, because its expected payment is based on the number of coin flips collected and, as before, one can always collect more coin flips in the forward direction (as going backwards necessitates visiting a node that's been visited before). In this case, the top player at $i + j$ would route up because the $i + k_i$-th player will have more coin flips (1.6ϵ) and is within the scope of the routing payment. Given this, it is strictly better for player i to route up instead of down, given knowledge of the values at A and B.

Consider now game G'''. By backwards induction, each player strictly prefers to route forward because doing so guarantees the largest payment along the way for any lookahead. The top player at $i + j$ will route along the middle path in equilibrium because he would receive $\epsilon + 1.7\epsilon$ from coin flips at the middle path of $i + k_i$ and $i + j + k_{i+j}$ vs. the 1.6ϵ $+\epsilon$ along the top path. In this case, player i would rather route down instead of up because it would collect 0.5ϵ more coin flips due to the 1.5ϵ at $i + k_i$ on the bottom path. However, since player i's best response routing decision should be the same for game G'' and G''', we have a contradiction. □

β	Dist.	$d = 4$			$d = 10$		
		MRSR	m=2	m=3	MRSR	m=2	m=3
.03	U	69	71	72	83	84	85
0.1	U	71	72	75	85	86	87
1.0	U	76	78	80	89	89	90
.03	S	80	87	104	150	183	227
0.1	S	88	109	146	181	226	259
1.0	S	120	155	183	227	258	278

Table 1: Comparison of routing performance after 10 steps on connected Watts-Strogatz graphs based on uniform (U) and skewed (S) coin flip distributions with fixed mean (5.5).

6. SIMULATION RESULTS

The equilibrium strategies induced by local routing rules can be considered to provide a heuristic algorithm for computing an optimal route over a network. We now demonstrate via simulations that routing decisions based on local rules can effectively aggregate information as a task is routed through the network.

We consider connected random graphs with 100 nodes and average degree $d \in \{4, 10\}$, generated using the Watts-Strogatz model [17]. By varying the re-wiring probability β, the model allows us to generate graphs that interpolate between a regular lattice ($\beta = 0$) and a $G(n, p)$ random graph ($\beta = 1$), with small-world networks emerging at intermediate values of β. We associate each node with a number of coin flips, which is drawn independently either discretely from U[1,10], or from a skewed distribution where the value is 1 with probability 0.9 and 46 with probability 0.1. Note that the distributions have equal mean (5.5), but that the skewed distribution more closely resembles a setting where there are few experts. For graphs generated in this manner, we simulate player strategies under local routing rules (MRSR, and m-hop with $m = 2$, $m = 3$) by computing local paths in the manner noted in Theorem 4, where revisited nodes are treated as having no value. As a baseline, we consider a random routing rule that routes to a random neighbor, and whenever possible, to a random neighbor who has yet to be assigned the task. Note that the expected performance of the baseline is bounded by 5.5 coin flips per round, as we would expect from randomly picking unvisited nodes in the graph.

Table 1 shows the average number of coin flips collected after 10 steps by players following local routing rules on graphs with varying β, average degree, and coin flip distribution over 100 trials (standard errors are small and hence not reported). We see that routing rules are particularly effective in cases where there are few experts (S), and when the graph has a sufficiently high connectivity (higher d and β) that there exist paths through which experts can be routed the task. But even in cases with uniformly distributed coin flips (U) and low average degree ($d = 4$), local routing rules collect significantly more coin flips than the upper bound of 55 we would expect from randomly choosing nodes, and that despite connectivity constraints the paths include many high valued nodes (recall the max per node is 10).

The difference in routing performance among local routing rules is rather small for uniformly distributed values, but is more significant when the distribution is skewed. In

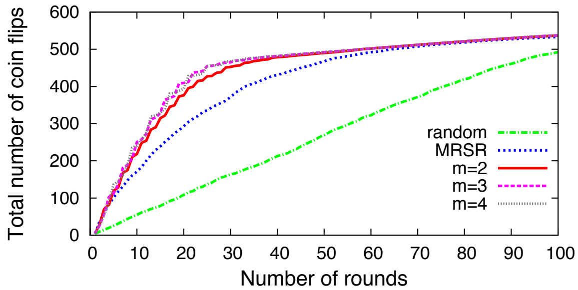

Figure 4: Comparison of routing performance among local routing rules for graphs with $\beta = 0.1$, $d = 10$, and skewed coin flip distributions. Values are averaged over 100 trials.

this case, effective routing may require finding short paths to experts who are not neighbors. That said, this difference shrinks for graphs with higher degree, as high-value nodes become more easily reachable (recall that as graphs approach cliques, myopic is optimal). Figure 4 shows the average number of coin flips collected by local routing rules as we progress through the routing game on graphs with $\beta = 0.1$, $d = 10$, and skewed coin flip distributions. We see that for $m \geq 2$ the performance under the routing rules are essentially the same, suggesting that we can sometimes achieve near-optimal performance globally with just two-hop local common knowledge. Note that for all local routing rules the rate of information aggregation eventually slows down, which denotes the point at which virtually all experts have been routed the task.

7. CONCLUSION

We consider the opportunity for incentivizing the joint refinement and routing of tasks among agents within a network, focusing on prediction tasks. We introduce and study routing scoring rules which, in equilibrium, support agents truthfully contributing information, and routing tasks based on simple computations that nevertheless lead to effective information aggregation. Future work on task routing for prediction tasks includes efforts to integrate additional information structures and study routing performance under specialized network topologies, consideration of intrinsic values for solving or routing, and introduction of communication or sensing mechanisms coupled with means of tracking costs for acquiring additional bits of signal. There are multiple opportunities to address task-level issues, and also organizational issues related to distributing streams of tasks in a manner that takes into account people's solving and routing abilities over a spectrum of tasks, as well as participants' changing levels of attention, motivation, and availability, and the corresponding need for balancing the load across participants. We envision numerous potential applications of methods for jointly solving and routing tasks and foresee an ongoing need to strike insightful balances between principled procedures and designs that rise from intuitions about practical implementations.

Acknowledgments

The problem was initially formulated at Microsoft Research during an internship there by Haoqi Zhang. Haoqi Zhang is funded by a NSF Graduate Research Fellowship. The research was also sponsored in part by the Army Research Laboratory under Cooperative Agreement Number W911NF-09-2-0053 and by National Science Foundation under Grant Number CCF-0953516. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

8. REFERENCES

[1] L. A. Adamic, R. M. Lukose, A. R. Puniyani, and B. A. Huberman. Search in power-law networks. *Phys. Rev. E*, 64:046135, Sep 2001.

[2] E. Arcaute, A. Kirsch, R. Kumar, D. Liben-Nowell, and S. Vassilvitskii. On threshold behavior in query incentive networks. In *Proc. EC '07*, 2007.

[3] J. E. Bickel. Some comparisons among quadratic, spherical, and logarithmic scoring rules. *Decision Analysis*, 4:49–65, June 2007.

[4] Y. Chen, D. M. Reeves, D. M. Pennock, R. D. Hanson, L. Fortnow, and R. Gonen. Bluffing and strategic reticence in prediction markets. In *WINE*, 2007.

[5] D. Dikshit and N. Yadati. Truthful and quality conscious query incentive networks. In *WINE*, 2009.

[6] P. S. Dodds, R. Muhamad, and D. J. Watts. An experimental study of search in global social networks. *Science*, 301(5634):827–829, 2003.

[7] J. Douceur and T. Moscibroda. Lottery trees: Motivational deployment of networked systems. In *SIGCOMM '07*, 2007.

[8] Y. Emek, R. Karidi, M. Tennenholtz, and A. Zohar. Mechanisms for multi-level marketing. In *Proc. EC '11*, pages 209–218, New York, NY, USA, 2011. ACM.

[9] R. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):3–15, February 2007.

[10] D. Horowitz and S. D. Kamvar. The anatomy of a large-scale social search engine. In *Proc. WWW '10*, pages 431–440, New York, NY, USA, 2010. ACM.

[11] J. Kleinberg. Complex networks and decentralized search algorithms. In *Proc. International Congress of Mathematicians*, 2006.

[12] J. Kleinberg and P. Raghavan. Query incentive networks. In *FOCS '05*, 2005.

[13] G. Pickard, W. Pan, I. Rahwan, M. Cebrian, R. Crane, A. Madan, and A. Pentland. Time-critical social mobilization. *Science*, 334(6055):509–512, 2011.

[14] R. Selten. Axiomatic characterization of the quadratic scoring rule. *Experimental Economics*, 1(1):43–61, June 1998.

[15] J. Travers and S. Milgram. An experimental study of the small world problem. *Sociometry*, 32:425–443, 1969.

[16] D. J. Watts, P. S. Dodds, and M. E. J. Newman. Identity and search in social networks. *Science*, 296(5571):1302–1305, 2002.

[17] D. J. Watts and S. H. Strogatz. Collective dynamics of "small-world" networks. *Nature*, 393(6684):440–442, June 1998.