Chapter 5

Task Routing

Engaging a crowd to tackle complex tasks relies not only on effective coordination, but on recruiting individuals with relevant expertise to join the problem-solving effort. One approach for bringing expertise to tasks is to pool knowledge about people's competencies and preferences and assign tasks in a centralized manner. Another approach is to rely on individuals in a system to select tasks themselves. Both approaches have flaws. In the former, a system may not know which individuals have the required expertise. In the latter, while individuals are often able to gauge their own expertise, they may not know which tasks best match their respective competencies.

In social networks, an individual's knowledge extends beyond their own expertise on tasks and topics to knowledge about the expertise of others. Members of a social network may know who among their friends can best answer a particular question or 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 people who would likely know experts. They may also be able to identify subsets of individuals who share a particular interest, among whom the requisite expertise is likely to exist.

We are interested in principles and methods for *task routing* that draw on the distributed intelligence of individuals across a social network. The idea is to harness the ability of people to contribute to a solution *and* route tasks to others who they believe can effectively solve and route. Task routing provides a paradigm for problem solving in which individuals in a crowd become engaged with tasks based on their peers' assessments of their expertise. On the task level, effective task routing aims to take advantage of people's knowledge about solving problems as well as people's knowledge about others' abilities to contribute. People make routing decisions in a peer-to-peer manner, and the system rewards participants for their contributions, both direct and indirect through routing. On the organizational level, task routing may provide a means for bringing tasks to individuals effectively, where people's routing decisions take into account not only an individual's expertise on the particular task, but also their ability to contribute as a router.

In this chapter, we focus on the special case in which the task is to obtain an accurate probability assessment about an uncertain event. The task is passed among individuals in 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 network structure and the amount of information held by everyone in the network, truthful reporting of posterior probability assessments and optimal routing can be obtained in a Perfect Bayesian Equilibrium. While this result is theoretically sound, optimal routing is NP-hard, which suggests that people may have difficulty computing routing decisions in practice. The common knowledge assumption is also unlikely to hold for large social networks, where each person's information about the competencies of others is limited to a local neighborhood (e.g., friends, and perhaps friends of friends).

To address these concerns, we consider designing incentive schemes for task routing problems where knowledge about the network structure and others' abilities is limited to an individual's local neighborhood. The main contribution is the introduction of a family of *local routing rules*, that isolate simple routing decisions in equilibrium under local knowledge about others' expertise and take advantage of such local knowledge to promote effective routing decisions. We achieve this by incentivizing participants to make routing decisions based on short, locally optimal paths that can be computed easily using local knowledge. In essence, *we design incentive schemes that explicitly enable equilibrium behavior for which the inference required of participants 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 best responses are invariant to knowledge outside of a local neighborhood. Simulation results demonstrate that equilibrium routing strategies based on local routing rules lead to effective information aggregation.

[1] This is analogous to the role of strategy-proofness in simplifying strategic problems facing agents in mechanism design [71].

5.1 Related Work

Methods for automated and manual routing of tasks have been employed in real world online networks. For example, question-answering services such as Aardvark [31] 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; a user can manually route questions to others, which enables the system to reach users outside its fund of knowledge about people and their expertise. Aardvark does not explicitly reward contributions, and instead relies on people's goodwill and social connections. In studying incentive mechanisms for task routing, we are exploring how properly rewarding participants for their contributions can help promote contributions to problem solving and routing more broadly.

Leveraging individuals' abilities to both solve and spread the word about the task was a key component of the winning team's strategy in the DARPA Red Balloon Challenge [73]. The task was to find large 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.[2] This mechanism aims to induce participants to broadcast the task to everyone they know, and is well-suited for one-off, high-stake scenarios such as search and rescue operations for which the benefit of reaching a large audience significantly outweighs the cost of people's atten-

[2]The interested reader may refer to Emek et al. [26], Douceur and Moscibroda [22], and Drucker and Fleischer [24] for related theoretical analysis, and related work on query incentive networks [48, 4, 19] that analyze games in which players split rewards to recruit others to answer a query.

tion. In contrast, the mechanisms in our work aim to leverage the expertise within a network by bringing to people's attention the tasks that they can best contribute to. These mechanisms are well-suited for efficiently processing a stream of tasks, without overloading people with information on every task.

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 [92, 21, 99, 47, 1]. In such work, the goal is to identify a *single* target node representing a particular individual. While this differs from our task routing problem, the results still provide 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 [29], which provide proper incentives for individuals participating in a prediction market to improve probability estimates by contributing additional information. The major difference between task routing and a prediction market is in who takes on the burden of identifying expertise. While a prediction market places 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.

5.2 Task Routing for Prediction Tasks

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 Boston Celtics win the NBA championship this year?" We consider discrete state spaces, and assume without loss of generality a binary state space, such that $\Omega = \{Y, N\}$.[3]

Consider a *routing 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 may be determined by the system or by the individual posting 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 prespecified number of rounds R, when a final assessment must be made. We assume players receiving the task are provided with a list of people who have participated so far, as well as the number of rounds that remain. Players are allowed to route to players who have participated thus far, but know that past participants may not have any additional information to contribute and may only be able to help with routing. *Our goal is to arrive at an accurate assessment after R rounds by designing 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.*

[3]For an event with more than two states, the task is to gather a probability vector with a likelihood assigned to each state. We can handle such events by using multi-class versions of proper scoring rules, and all of our results extend straightforwardly.

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.[4] We assume the realization of each player's signal is private, and make different assumptions about the knowledge of a 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

[4]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.

length, as long as previous updates were done truthfully [14]. The posterior incorporates, and sufficiently summarizes, all information collected thus far. To perform updates, players need to only know the signal distribution with respect to their own signal, which we assume is known to all players. 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.

5.3 Routing Scoring Rules

With rational, self-interested players who have no intrinsic value (or cost) for solving or routing a particular task, ensuring effective task routing requires mechanisms that will incentivize players to both truthfully update posterior probabilities and 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* [83] are mechanisms that strictly incentivize a forecaster to truthfully reveal his subjective probability of an event, typically under the assumption that participants are risk neutral. The outcome of the event is assumed to be 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 well-known strictly proper scoring rules include the logarithmic and spherical scoring rules. Any strictly proper scoring rule can be scaled or normalized via linear transformations to form another strictly proper scoring rule [7].

Market scoring rules [29] 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 player's report. 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.

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

Definition 5.1. *A* **routing scoring rule** *defines a sequence of positive integers* k_1, *...,* k_{R-1}*, which rewards players* $i \in \{1, \ldots, R-1\}$ *on the routing path:*

$$(1-\alpha)s_i + \alpha s_{i+k_i} - s_{i-1} \tag{5.1}$$

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 a combination of his report and the report of the player who receives the task k_i steps after him, relative 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 $s_R - s_{R-1}$. Players who participate multiple times within a routing game are paid based on the routing scoring rule the first time they receive the task, and paid by the market scoring rule in any subsequent interactions.[5]

Intuitively, routing scoring rules reward players who are experts as well as players who are knowledgeable about the expertise of other players. We introduce 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 adjacent player who can most accurately refine the probability assessment.

Lemma 5.1. *Consider a routing game in which each player participates at most once. The total payment from the system 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. The expression is familiar from market scoring rules, containing just an additional term due to routing payments.

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

[5] For the local knowledge settings we consider later in the chapter, this avoids situations in which a player may try to hold on to a task by making suboptimal routing decisions that lead to their being routed the task again, with the intent of earning multiple routing payments beyond the first.

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, judged at the end of the game.

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

Lemma 5.2. *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 the number of bits l_j for players $i < j \leq k$ on S' and S''. Let $E^i_S[s_k]$ denote player i's expectation, taken immediately after his own report, 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) components: (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 component 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.

5.4 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 network structure and the number of coin flips l_i observed by each player i is common knowledge. Note the actual signal realizations are still assumed private.

5.4.1 Clique Topology

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

Theorem 5.1. *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 $i+1 \in \operatorname{argmax}_{m \in S_{>i}} l_m$, 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. For any player i, 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. Note that for (a), since s_i is based on a strictly proper scoring rule, truthful reporting maximizes the expectation of s_i. For (b), 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 5.2. □

5.4.2 General Networks

We now consider routing games on general networks with missing edges; e.g., only managers can route tasks between teams 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 5.1. *Consider the routing graph $G = (V, E)$, in which 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.*

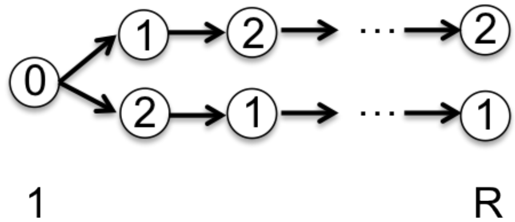

Figure 5.1: A routing game for which myopic routing (along the bottom path) is suboptimal. Numbers in nodes are the number of coin flips held by each player.

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

Immediately, we see that myopic routing will not always find the optimal solution to Problem 5.1, as routing to the neighbor with the most coin flips does not consider the effect this can have on future routing decisions, and can now convey an opportunity cost. Figure 5.1 illustrates an example in which myopic routing would lead player 1 and all subsequent players to route along the bottom path, while the optimal solution requires players to route along the top path.

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

Lemma 5.3. *Problem 5.1 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.

But intractability is not the only difficulty we face. 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 5.2. *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 5.1 for which $k = R - i$, $o = i$, and $w_m = l_m$ if $m \in S_{>i}$ and $w_m = 0$ otherwise. Under the path-rewarding routing scoring rule, it is a PBE for each player i to truthfully update the posterior probability and route the task to the next player in the path provided by Q_i, with the belief that all other players follow this strategy.*[6]

Proof. (sketch) Using similar arguments as in the proof sketch for Theorem 5.1, we show that no player wishes to deviate from the equilibrium strategy, given the belief that all other players report truthfully. For any player i, we first show that player i should honestly update the posterior beliefs by establishing that (a) truthful reporting maximizes s_i, and (b) for any subsequent sequence of players $i+1, \ldots, R$ who may be routed the task, truthful reporting by player i maximizes the score s_R at the end. For (a), since s_i is based on a strictly proper scoring rule, truthful reporting maximizes the expectation of s_i. For (b), the expected score of s_R (from the perspective of player i) is strictly greater when player i reports honestly because s_R is based on a strictly proper scoring rule.

[6]In this setting, a player who participates multiple times does not receive, nor require, any incentives for routing beyond the first time. This is because routing along an optimal path is required for maximizing the expected score at the end of the game, which is the basis for a player's (first time) routing payment under the path-rewarding routing scoring rule.

It is left to show that player i maximizes s_R by routing to the next player in the path provided by Q_i; this follows from Lemma 5.2. □

Since PRSR rewards each participant's routing decision based on the final score, it is in each participant'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:

Theorem 5.3. *The set of PBE identified in Theorem 5.2 (corresponding to possible ties in the solution to problem 5.1) 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 participants' 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 5.2). □

5.5 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. An individual will not necessarily know everyone else, and may only have limited 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 other's expertise within a local neighborhood of a particular size, defined by the number of hops between participants. For example, all friends of a particular person are aware of his expertise (one hop). Friends of his friends may also be aware (two hops).

Definition 5.2. *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 that 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 from having more information.

Given that a player may only have m-hop local common knowledge, let's consider the problem facing such a player when deciding how to route to maximize the final prediction quality after R steps. Routing optimally may require the 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 held by different people in the network. Furthermore, optimal routing requires a player to make inferences about the values 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 will likely be fragile because it requires players to adopt priors over other players' beliefs.

An attempt to avoid such issues may suggest rewarding 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 would still have to consider the routing decisions of players outside his locality because maximizing his payoff requires considering the routing decisions of the chain of players within his locality. For example, consider the two-step routing rule (see bottom of Figure 5.2). 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 paid for the score two steps forward (from him), 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, in order to just compute the expected score after two steps.

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

Definition 5.3. *A player i in a routing game adopts a* **m-local strategy** *if his routing decision depends only on m-hop local common knowledge and is invariant to any beliefs the player might have about players outside of his own locality.*

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

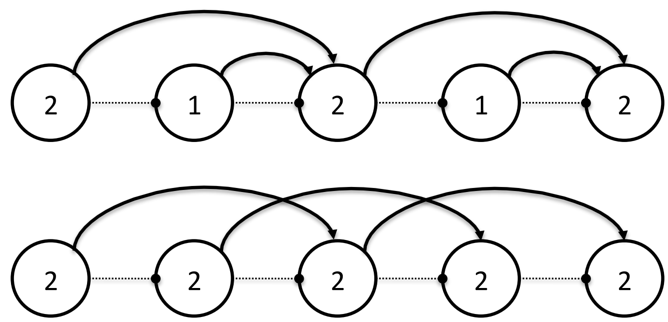

Figure 5.2: Illustration of the 2-1-2-1 and 2-step routing rules. Arrows depict dependencies in routing payments.

Definition 5.4. *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 5.2), which can be computed with local common knowledge. As even-numbered 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-numbered player knows the number of coin flips that can be collected from the next even-numbered player and the next odd-numbered player that is routed the task, he can compute the best local path without regard to routing decisions beyond his locality. 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 5.5. *The* $\boldsymbol{m}$**-hop routing rule** *is a routing scoring rule which sets* $k_i = \min[m - (i-1) \bmod m, R - i]$.

The m-hop routing rule supports the following equilibrium behavior:

Theorem 5.4. *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 5.1 for which $k = \min[m - (i - 1) \bmod m, R - i]$, $o = i$, and $w_j = l_j$ if $j \in S_{>i}$ and $w_j = 0$ otherwise. Under the m-hop routing rule, it is a PBE for each player i to truthfully update the posterior probability and 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 5.1, 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 5.2 proves the point, and the theorem. □

Unlike in the common knowledge setting under the path-rewarding routing scoring rule, this equilibrium under the m-hop routing rule may not be unique. For a player routing more than once, after the first time, the player is weakly indifferent among all routing decisions because his payment reduces to the market scoring rule for subsequent routing opportunities. Such a player need not route along a locally optimal path in making subsequent routing decisions and this can affect the equilibrium behavior of other players who may route the task back to this player. If we wish to ensure that routing along a locally optimal path is a unique equilibrium, we can modify the routing game slightly to prevent players from routing to other players who have already participated in the game.[7]

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. It turns out that 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 the next two players are each paid myopically. Note that the first player here can still compute his optimal routing decision using only local common knowledge by computing the routing decisions of others in his locality via backwards induction. We can thus characterize the entire family of local routing rules:

Definition 5.6. *Given m-hop local common knowledge, the family of* **m-local routing rules** *consists of routing scoring rules defined with parameters* $k_1, \ldots, k_{R-1}$*, that satisfy* $k_{i+j} + j \leq m$ *for all* i *and* $0 \leq j < k_i$.

Generally, we can refer to these as *local routing rules*, dispensing with the m when this detail is unimportant. The 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 this condition for $m = 2$ because

[7] We can modify the routing game so that in cases when a player has no one to route to, no routing payments will be assigned. The task is returned to the system which will randomly select a new participant. Since players cannot participate twice in this modified game, uniqueness of the equilibrium stated for the m-hop routing rule in Theorem 5.4, and more generally for local routing rules in Theorem 5.5, can be recovered without requiring further assumptions. The argument is similar to that in the proof of Theorem 5.3.

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 the condition, because for all $i < R-2$, $k_{i+1}+1 = 3 > m$. Note that the m-hop routing rule satisfies the condition, since k_i is set such that $k_{i+j} + j = m$ for all appropriate i and j in Definition 5.6.

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

Theorem 5.5. *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 m-local routing rule, consider the following dynamic program:*

$$\begin{aligned} V(n_{j+1}, \ldots, n_{j+k_j} | n_1, \ldots, n_j) = \max_{j+1,\ldots,j+k_{j+1}} [\sum_{b=1}^{k_{j+1}} w_{j+b} \\ + V(n_{j+k_{j+1}+1}, \ldots, n_{j+k_j} | n_1, \ldots, n_{j+k_{j+1}})] \end{aligned} \tag{5.2}$$

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

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 participants 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 a player's own assessment and the assessments collected from those routed the task via the routing payment. Second, as

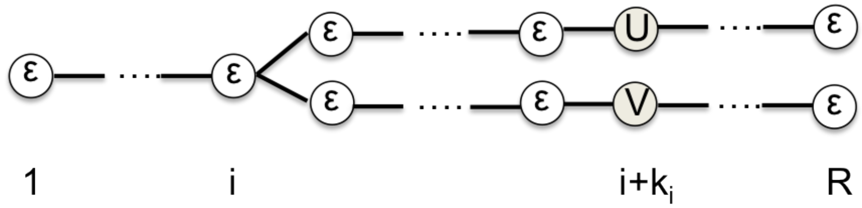

Figure 5.3: Routing game construction for the $j = 0$ case.

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 definition of local routing rules $k_i \leq m$, players follow m-local strategies. That is, 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 5.2 proves the point, and the theorem. □

Theorem 5.6. *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 is not a local routing rule, 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 shown in Figure 5.3.

Based on the construction, consider two routing games G and G'. 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 the condition for local routing rules 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 participants 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. This is because a player's expected payment is based on the number of coin flips collected and one can always collect more coin flips in the forward direction (for any player, going backwards would necessitate visiting a node that's been visited before with 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 shown in Figure 5.4.

Based on the construction, consider two routing games G'' and G'''. 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 the condition for local routing

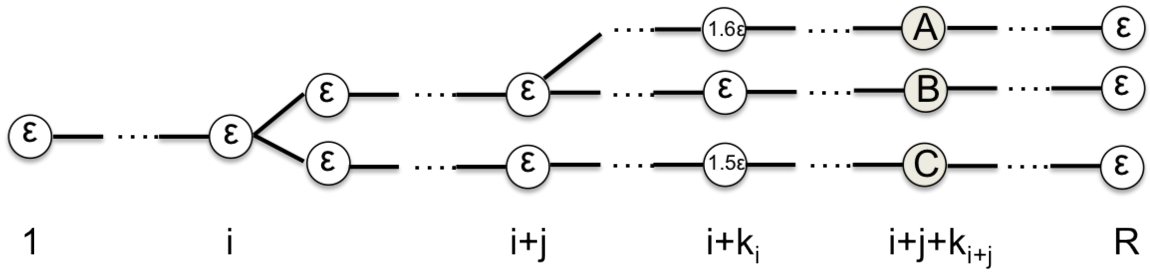

Figure 5.4: Routing game construction for the $j > 0$ case.

rules, 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 players 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 all times, regardless of the lookahead induced by their routing payment. This is because a player's 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 would have more coin flips (1.6ϵ) and is within the scope of the routing payment. Given knowledge of the values at A and B, it is thus strictly better for player i to route up in G''.

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 $i + j + k_{i+j}$ versus the $1.6\epsilon + \epsilon$ 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. □

5.6 Simulations and Results

The equilibrium strategies induced by local routing rules can be viewed as providing 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 [100]. 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. Coin flips are 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. The two distributions have equal mean (5.5), but 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 5.4,

		$d = 4$			$d = 10$		
β	Dist.	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 5.1: Comparison of routing performance based on the average number of coin flips collected after 10 steps. Values represent averages over 100 trials. We considered connected Watts-Strogatz graphs based on uniform (U) and skewed (S) coin flip distributions with fixed mean (5.5). In all cases, routing based on local routing rules collected significantly more coin flips than the 55 coin flips (upper bound) we would expect to collect from routing randomly.

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 5.1 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) such that paths exist 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

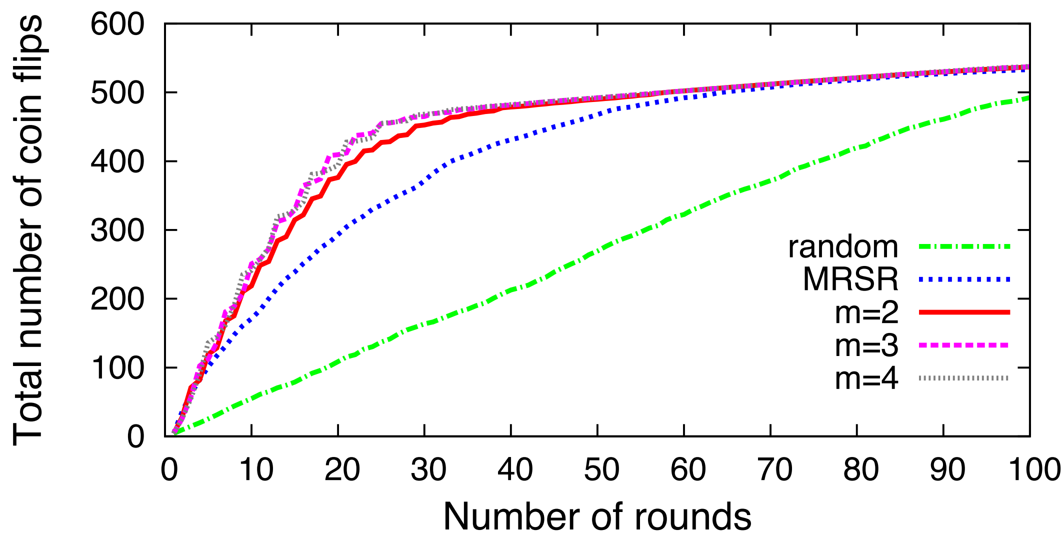

Figure 5.5: Comparison of routing performance based on the average number of coin flips collected for graphs with $\beta = 0.1$, $d = 10$, and skewed coin flip distributions. Values represent averages over 100 trials. Routing based on local routing rules collected significantly more coin flips over fewer rounds than routing based on the random routing rule.

from randomly choosing nodes. Despite connectivity constraints, paths still included 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 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 reachable (recall that as graphs approach cliques, myopic is optimal).

Figure 5.5 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 routing based on local routing rules collected significantly more information over fewer rounds than routing based on the random routing rule. For $m \geq 2$, the performance under the local routing rules are essentially the same, suggesting that we can sometimes achieve near-optimal performance globally with just two-hop local common knowledge.

With the random routing rule, we see that the rate of information aggregation stays nearly constant throughout the routing game. Since the rule routes to new players whenever possible, this suggests that the graph is well-connected and that new players can often be routed the task even later in the game when many players have already participated. With local routing rules, we see that the rate of information aggregation eventually slows down, which denotes the point at which virtually all experts have been routed the task.

5.7 Discussion

We consider the opportunity for incentivizing the joint refinement and routing of tasks among people within a network, focusing on prediction tasks. We introduce and study local routing rules which, in equilibrium, support people truthfully contributing information and routing tasks based on simple computations that nevertheless lead to effective information aggregation.

In our analysis, we have assumed that bits of signal are conditionally independent. But in some settings, players' signals may be *conditionally dependent*, and accurate predictions may depend on collecting the complementary information held by different players. In this setting, our theoretical results continue to hold with small modifications. First, it is no longer sufficient to maintain a posterior estimate. Instead, we need to explicitly keep track of the information contained in players' signals. Second, we need to restrict players from participating more than once, or alternatively, from being paid for their information beyond the first time. This prevents the type of incentive issues that may occur in prediction markets, in which participants with conditionally dependent signals may be better off withholding some information until complementary information has been reported to the market [14].

While local routing rules enable equilibrium behavior for which the inference required of participants is tractable, these rules still assume that participants are rational in that they maximize their expected payoff. In practice, participants can make mistakes and route suboptimally. But even so, local routing rules may provide for a robust design in which participants are incentivized towards making good decisions even if their decisions are not optimal. Since local routing rules are based on strictly proper scoring rules, which in our setting are *accuracy-rewarding* [51], more accurate predictions will lead to strictly higher payoffs. Furthermore, since the equilibrium is constructed within local paths, any "mistakes" also remain local, and do not affect the routing decisions of later participants outside of local reach.

In crafting local routing rules, we demonstrated a means for designing incentives that explicitly enable players to make simple computations in equilibrium. The key idea is to ensure that players need only make decisions based on information they are guaranteed to have. This requires that players' routing payments are localized and that any chains of reasoning are limited to within local neighborhoods. We believe this idea generalizes beyond prediction tasks and can enable effective solving and routing over social networks in a variety of settings.

There are many possible directions for future work on task routing. One direction is to study routing performance under specialized network topologies and knowledge distributions. Another direction is to extend our models to consider the intrinsic value and cost for solving or routing. In this direction, we are also interested in introducing communication or sensing mechanisms coupled with means of tracking costs of acquiring information, in order to take into account and study the process through which individuals make contributions.

We are interested in developing general principles and methods for solving and routing with humans and machines, and in considering other types of meta-knowledge participants may have about the expertise of others in a social network. In addition to multiple opportunities to address task-level issues, there are also opportunities to address 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. We elaborate on this direction in Chapter 9.