

Lattice Games and the Economics of Aggregators

Patrick R. Jordan
 Yahoo! Labs
 Santa Clara, CA 95054
 prjordan@yahoo-inc.com

Uri Nadav
 Yahoo! Research
 Santa Clara, CA 95054
 urinadav@yahoo-inc.com

Kunal Punera †
 RelateIQ
 Mountain View, CA 94041
 kunal.punera@utexas.edu

Andrzej Skrzypacz ‡
 Stanford University
 Stanford, CA 94305
 andy@gsb.stanford.edu

George Varghese ‡
 UC San Diego
 La Jolla, CA 92093
 varghese@cs.ucsd.edu

ABSTRACT

We model the strategic decisions of web sites in content markets, where sites may reduce user search cost by aggregating content. Example aggregations include political news, technology, and other niche-topic websites. We model this market scenario as an extensive form game of complete information, where sites choose a set of content to aggregate and users associate with sites that are nearest to their interests.

Thus, our scenario is a location game in which sites choose to aggregate content at a certain point in user-preference space, and our choice of distance metric, Jacquard distance, induces a lattice structure on the game. We provide two variants of this scenario: one where users associate with the first site to enter amongst sites of equal distances, and a second where users choose uniformly between sites at equal distances. We show that subgame perfect Nash equilibria exist for both games. While it appears to be computationally hard to compute equilibria in both games, we show a polynomial-time satisficing strategy called Frontier Descent for the first game. A satisficing strategy is not a best response, but ensures that earlier sites will have positive profits, assuming all subsequent sites also have positive profits. By contrast, we show that the second game has no satisficing solution.

Categories and Subject Descriptors: J.4 [Computer Applications]: Social & Behavioral Sciences

General Terms: Economics

Keywords: Content aggregators, Game theory

1. INTRODUCTION

The process by which users find and consume digital content, browsing for short, has transitioned through several phases since the advent of the Internet. In early periods, two styles of browsing were prevalent: *search engines* (culminating in sites such as Google) and *directories* (Jerry and David's Guide which culminated in Yahoo). In terms of user interaction, search engines offered low search cost, given that the user knew *a priori* the specific topic they were interested in. Comparatively, directories allowed users to start at a very general area of interest and gradually refine content until a suitable set was found. Hence, the burden of search was left to the user, but they did not need to have a specific topic of interest in mind before browsing. Portals, websites consisting of links or hosted content specific to a topic, emerged as a middle ground be-

tween search engines and directories. The idea of a *favorite portal* took hold, and sites like AOL and Yahoo attempted to cover the set of user interests with a small set of portals each devoted to distinct, but high-level, topics. Conceptually, this middle ground would allow users to visit a small set of websites to consume their fill of content, with relatively low search costs.

Recently, strategic behavior of some providers and unabated growth of online content have started undermining the standard search-for-topic and browse-through-favorite-portal paradigms. First, search engine optimization and spam make it hard to find what one wants when searching for some topics, such as cameras. Second, along with inferior content proliferating, well-crafted niche sites keep appearing—for instance, StarFall (early reading) or MacRumors (an aggregator of Apple-related news). The illusion that one can satisfy oneself with a small set of portals is becoming harder to sustain; yet, user attention and clicks remain the primary, limited resource used to discern the gems among the garbage.

This qualitative argument may explain why web site aggregators have become such a powerful and rising phenomenon. Rather than read thousands of individual technical sites or more general web sites such as Yahoo and AOL, many software professionals get their technical news from some combination of TechCrunch, Slashdot, Ars Technica, and AllThingsD (and the like). There are subtle differences between the offerings: AllThingsD appears to specialize more in technical gossip (e.g. CEOs fired), whereas Slashdot seems to specialize in technical ideas. There is further specialization: for example, AppleInsider combs the web for Apple news and Mac software updates. Such catering to specific tastes is widespread: for example, while many people consume news with broad sites such as Yahoo! News and CNN News, a large constituency prefers to consume news with a political slant, at say Huffington Post (liberal) or Drudge Report (conservative). Aggregators exist for photography (Mashfot), Nintendo fans (WiiNintendo), designers (MostInspired). If one believes in the phenomena of the long-tail that explains the success of Amazon and Netflix, then it seems reasonable to posit that increasingly specialized aggregators (e.g., Indian Immigrant children) will keep appearing.

What incentives do aggregators have? Why do aggregators enter some spaces and not others? At what point does specialization lead to losses? In this paper, we attempt to answer these questions from a game-theoretic perspective. We develop a model where aggregators' and users' strategies are modeled as subsets of a con-

Copyright is held by the International World Wide Web Conference Committee (IW3C2). Distribution of these papers is limited to classroom use, and personal use by others.
 WWW 2012, April 16–20, 2012, Lyon, France.
 ACM 978-1-4503-1229-5/12/04.

[†]This work was done when the author was a Senior Research Scientist at Yahoo! Research.

[‡]This work was done when the author was visiting Yahoo! Research.

tent universe. A *user set* models user preferences, describing the content that a user at a given point in time is interested in. An *aggregator set* models the content the aggregator hosts. Users have a simple strategy: they pick the *closest* aggregator (which we model as the Jacquard distance between their user set and the aggregator set). Aggregators, on the other hand, seek to maximize the number of users that select them. We consider a sequential game in which aggregators enter serially, if they can make profit and position themselves on nodes in the content lattice. Aggregators incur a fixed cost F upon entry. After all entry decisions are made, each user selects an aggregator.

The tension in the model is that specialization can decrease Jacquard distance and cause some users to switch, but generalization can capture a larger set because a user that cannot find the perfect aggregator (one with a set exactly equal to the user's taste) prefers to go to an aggregator that has k additional content pieces than one that is missing k pieces (a property of our distance measure that we find realistic).

The questions we ask include: When and where should an aggregator enter? In a game with infinite number of potential entrants, is there a subgame perfect Nash equilibrium? Is it in pure strategies? Under what conditions is there entry in equilibrium? What is the maximum/minimum number of equilibrium entrants? Are there reasonably computable strategies that an aggregator can use to decide when and where to enter?

2. RELATED WORK

The phenomenon of specialization and strategic entry is certainly present in other markets. The seminal early papers on location games in one-dimensional taste space are Hotelling [6] and Salop [9]. These models have been extended in many ways, see for example Economides [5], Caplin and Nalebuff [3], Ansari, Economides and Steckel [2] among many others. While some of these extensions consider multiple product characteristics, there are two important differences from our model. First, the location space is very different—to the best of our knowledge, in all these papers location is chosen on a subset of \mathbb{R}^N . By contrast, we model content aggregators as choosing subsets of the universe of available content, introducing a lattice structure on the location space and a different distance metric that is associated with it. Second, we do not have prices (which play an important role in the models of product differentiation) because most existing aggregators offer free access to users¹ (this second assumption makes our model closer to a model of a strategic platform choice in political economy models, but these also differ from our model with respect to the location space and they typically look at a winner takes all competition).

Our model is also related to issues of product (or store) design. If there is a universe of all features that a product may have, then each producer designs its product expecting future competition with other producers (who make their own designs). That problem is often modeled as a quality choice problem (starting with Shaked and Sutton [11]), but that makes it similar to the models of location choice. We are not aware of any models of product design that would treat the design as a selection of features and/or consumer preferences to be characterized by Jacquard distance, instead of the more usual Euclidean distance. Additionally, for our model to be applied to

¹Negative prices to attract users are impractical for obvious reasons, but as long as advertising is perceived by users as negative utility, aggregators could compete in the amount of advertising their sites have. Yet, it is possible that at least for some small level of advertising users do not mind advertising and competition between aggregators could drive the advertising to that level. For simplicity we abstract away from competition in the amount of advertising.

product design, it would be natural to allow for price competition and for the costs of the product to depend on the number of features. In contrast, it appears to us that the costs of aggregation of online content increase much less with the amount of aggregated content, hence we propose to assume these effects away.

Finally, another set of related literature is on bundling. These papers often focus on whether a particular producer should offer its products separately or to bundle them (for example, sell individual-game tickets separately, or sell them as season tickets²). Such considerations may be relevant for publishers that introduce paywalls (and then allow users to purchase subscription to multiple websites or purchase them a la carte), as well as to questions of firms managing multiple aggregators (a recent strategy of AOL). We leave the first topic for future research and provide partial discussion of the case where one firm can enter with multiple aggregators.

In summary, even though there is a vast literature in economics (and political economics) on entry decisions with differentiated products (differentiated by taste, location, features and bundling), we believe that our (new) model of entry on a lattice has several advantages over the existing models:

1. We believe Euclidean distance is a poor model (in any dimension) for content disparities between users and web sites. In some simple cases (conservative versus liberal), a Euclidean may be appropriate, but the nuances of taste such as TechCrunch versus MacScour, Yahoo News versus Huffington Post seems harder to capture. We believe Jacquard distance is a better model.
2. We model aggregators as having a fixed constant cost regardless of the size of content they host. This seems a reasonable assumption for virtual goods, such as content, but is unreasonable for physical goods, such as furniture, where there is an incremental stocking cost per SKU.
3. Standard economic models also add the notion of prices and competition for prices but for the most part most aggregation sites today are free and are paid for by advertising. The few sites that have attempted pricing (e.g., NYTimes) have unclear outcomes today. In the web world, aggregators attempt to capture users (and hence ad revenues) by tailoring content and not by lowering prices or reducing advertising below certain level needed for ease of use (with the notable exceptions of websites like Wikipedia that are not-for-profit).

Beyond a different model, some of our results also introduce a new way of analyzing the problem. One way of analyzing such a game is to characterize subgame perfect Nash equilibria (SPNE). While we accomplish this with stylized examples, providing many properties of SPNE outcomes is hard; and even computing SPNE seems to be a difficult problem computationally. Therefore, after we establish some basic results about SPNE, we move to finding a *satisficing strategy*: a strategy that, given the entry decisions of existing firms, a firm can take to guarantee itself non-negative payoffs under a minimal assumption that future entrants will take only strategies that would result in non-negative payoffs for themselves. That strategy is neither a best-response (required for equilibrium), nor it is a max-min strategy (which would allow the following entrants to lose money), but it is somewhat in between: we allow the future entrants to take adversarial actions, but constrain them not

²That literature starts with Stigler [13] and includes, among many others, Schmalensee [10], McAfee, McMillan and Whinston [7] and Nalebuff [8]. The last paper studies how bundling can be used to deter entry.

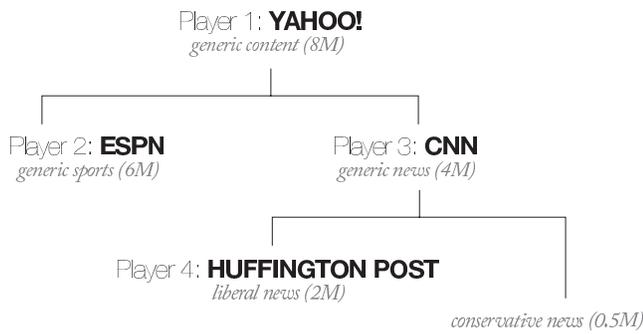


Figure 1: Example to motivate the lattice model

to lose money themselves. We prove for one of the variants of the model (consumers choose earlier entrants in the case of a tie) such a strategy exists and can be found by a simple greedy algorithm. In that model, if all entrants follow that algorithmic strategy, the final outcome is a limited entry with varied types of aggregators that all make positive profit. We show that this strategy is not necessarily a part of a SPNE (i.e., it is not a best response). Finally, we show that if ties are broken symmetrically, this algorithm does not yield a satisficing strategy.

3. MOTIVATION

The goal of our model is to explain the evolution of content consumption. Starting from impersonalized portals, which tried to satisfy every user with all possible pieces of content, there has been an evolution to another extreme (search engines) that are very granular in topic. However, aggregators provide a middle ground where users do not have to explicitly know what they want and yet do not have to perform the manual search inherent in portals. We would like to understand how these aggregators choose the set of content they provide. To do so we will model user preference as *content taste sets*, the entrance of aggregators into the market as an *extensive form game*, and the resulting *payoffs* that accrue to the aggregators as their incentive to enter the market.

Taste Sets.

To motivate the model, consider Figure 1. The figure represents sets of user tastes. For example, it posits a large set of 8 Million (8M) users who are interested in generic content. It also assumes that there are 4 million users interested in generic news, 6 million in generic sports. Finally, there are 2 Million users interested in liberal news and 0.5 million interested in conservative news. We draw an arrow from a more generic set of user tastes to a more specific set of user tastes so the sets form a lattice—the primary object we use to play our aggregator games.

Note, however, that the value of a *taste set* need not equal the sum of the values of all its subsets in the lattice. For example, in our example there are 4M users who are happy to view generic news, but there are 2M users who are only interested in news with a liberal slant and 0.5M who are interested in news with a conservative slant. Thus 4 is not equal to 2 + 0.5. This makes sense because the set of users interested in generic news can be interpreted as the set of users who are equally happy to consume both conservative or liberal news. This is completely independent from the set of users who overwhelmingly prefer liberal news.

One can assume that these content taste sets could be measured by surveys and are common knowledge to all players (website providers) who choose to cater to certain taste sets. Let us also assume that a

web site must have 1 million users to break even. We model that as a fixed cost of 1M assuming the currency is users.

User Behavior.

Assume that initially player 1 in the game (e.g., Yahoo) may choose to provide a generic content site. At that point, if Yahoo is the only site, all users will (in our model) flock to Yahoo (because there are no competitors yet) so that Yahoo gets 20.5 Million users and a payoff of $20.5M - 1M = 19.5$. However, this could encourage a second player (say ESPN) to start a sports web site. Now the users only interested in sports will switch to ESPN, and so ESPN garners 6M users and Yahoo loses those users. We model this mathematically by saying that the Jacquard distance of the Generic Sports users is smaller to ESPN than to the generic content site, and that users unilaterally switch to web sites have the smallest Jacquard distance to their own taste set. Recall that the Jacquard Distance between two sets is $1 - R$, where the resemblance R is the size of intersection divided by the size of the union of the two sets.

Intuitively, this makes sense because large Jacquard distances imply sites with a large amount of content that a user does not want; in some sense, we are modeling user unhappiness by a large Jacquard distance. Users thus maximize their utility by picking aggregators with the smallest Jacquard distance from their own taste set. We believe that Jacquard distance also makes sense for the standard economics of specialization in the physical world. For instance, users prefer to go to a furniture store to buy furniture rather than to a general purpose store such as Walmart. Standard economic literature uses Euclidean distance and location games to model this phenomena; one of our contributions is suggesting that Jacquard distance is a more accurate model to capture the nuances of user tastes.

The Extensive Form Game.

Continuing with the extensive form game of Figure 1, a player 3 (say CNN) may decide to enter the news market and claim 4M users away from Yahoo. Further specialization is possible because there is a still attractive market of 2 million liberal news users. Thus a Player 4 (say Huffington Post) may enter which steals away 2M users from Player 3. At this point, there is no incentive for a fifth player to enter because the conservative news market does not provide enough revenue (0.5M) to break even after a 1M fixed cost. If we chose to model Generic Sports as having further subcategories (e.g., Baseball, Football, Basketball) with sufficiently large revenue, then further players (e.g., MLB News, NBA News) could also have incentives to enter the game if the number of sports fans in each category are sufficiently large.

We model the situation in Figure 1 more abstractly in Figure 2. We represent by the string A the set of users with liberal news taste, by B the users with a conservative news bent, and by AB the set of users who are indifferent to the slant and can thus consume conservative and liberal news with equal relish. Similarly, we represent by C the set of users who like sports. ABC then represents the set of users who like generic content: these are users who are happy to consume sports and news of any kind.

Of course, generic content clearly includes other categories as well such as Entertainment and Technology. These can be modeled by extending our alphabet of characters but we have chosen not to do so in order to keep our example as simple as possible. There are also clearly other possible lattice nodes not shown in Figure 1, such as users who like either conservative news or sports but not liberal news (set BC). In all subsequent lattice diagrams, assume that any lattice nodes not explicitly shown have zero users.

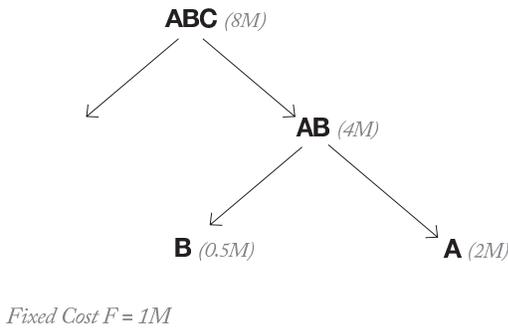


Figure 2: The lattice structure for the motivating example in Figure 1

Define the Jacquard Distance $JD(X, Y)$ between two strings X and Y as $1 - R(X, Y)$. The resemblance $R(X, Y)$ between two strings X and Y is defined as $I(X, Y)/U(X, Y)$. $I(X, Y)$ is the size of the intersection between the set of characters in X and the set of characters in Y . $U(X, Y)$ correspondingly represents the size of the union between the set of characters in X and the set of characters in Y .

Based on these definitions, it is easy to see that $JD(C, C) = 0$ while $JD(ABC, C) = 2/3$. In other words, users with a taste set of generic sports will prefer a sports site like ESPN (in our model), to a generic content site such as ABC which has been taken by say Yahoo.

Attributing Revenue.

We also have to make a major modeling choice as to what to do when two sites have the *same* Jacquard distance to a user choice set. There are two simple possibilities. The first is what we call FMT (First Movers Take Ties). In the FMT game, the site which comes first in time (recall we are playing a sequential game) wins all user choice sets of equal Jacquard distance. On the other hand, a more standard assumption in economic theory is the ET game; in the ET game, if there are ties the revenue is shared equally.

For example, in Figure 1 suppose that another liberal news site enters the fray. In the ET model, this is reasonable because a new site competing for the liberal news taste set can get a million users and break even. This is not true in the FMT model; the FMT game models situations where users do not switch (because of say inertia) if there is a new site that caters to exactly the same tastes. The assumption is that if the new site has exactly the same set of content, users prefer to stay with the existing site. Of course, the truth is more nuanced. Users may switch to the new site despite inertia because the voice of the site is subtly different, something that would be hard to model. We believe the real truth is somewhere between both models. Thus we study both models in the sequel.

Modeling Subtleties.

We do not explicitly model the following phenomena.

- *Variable costs for players:* In reality, the cost of a website varies depending on the number of users it serves if simply in terms of the costs of servers and electricity. However, this can easily be modeled by simply subtracting the variable cost from the fixed revenue of a node before placing it in the lattice.
- *Multiple moves by a player:* We allow a player to make only 1 move (to enter or not) and not to make further moves. To

some extent this models the fact that completely changing a web site to fit a new set of tastes is tantamount to completely restarting as a new player with a new fixed cost. Imagine the difficulty of retooling The Huffington Post to become a competitor to TechCrunch!

- *Users choosing multiple aggregators:* In practice users do not visit just one site that is closest to their interests but a small set of sites, while possibly minimizing a browsing budget. Modeling this seems very hard because picking a union of taste sets that minimize Jacquard distance seems akin to set cover which is computationally hard. As we will see, even with a simple single choice of web site, the games are structurally complex and appear to have hard to compute equilibria. Thus it makes sense to start with the simplest model and add complexity later. Further, instead of modeling a user as choosing B taste sets, we could alternately model an individual user as a probabilistic agent that chooses different taste sets with defined probabilities. For example, a user at any point in time may be in the mood for Sports with probability 0.8 and for news with probability 0.2. In that case, our simple model applies to some extent if the revenue numbers attached to the taste sets are interpreted as the expected number of users.
- *Sparse Lattices:* In practice, the vast majority of combinations of user taste sets will not be known and will be impossible to survey. Thus, practical models will have most of their weight(revenue) in the leaves of the lattice. We leave specializing our results to such “sparse” lattices as future work. For this paper, they can be modeled as nodes with zero revenue.

4. FORMAL MODEL

We now proceed formally. We have a countably infinite set of potential entrants I (e.g., Players like Yahoo, ESPN etc) with typical element $i \in I$. We have a lattice L of subsets of an underlying Content set C with a typical element $S \subseteq C$. We assume that every subset S in the lattice has a value $v(S)$ which represents the payoff for capturing users with that set of tastes; the values of $v(S)$ are weakly positive. For convenience, we will find it useful to define the descendant revenue $V(S)$ of a lattice node as the sum of the revenues of S and all descendants of S in the lattice. For example, in Figure 2, $v(ABC) = 8M$, the number of users who are interested in generic content but $V(ABC) = 20.5M$.

Player i at time i observes the history of the game and decides whether to enter and where to enter. The location of entry (player i action) is a node (a taste set) in the lattice, and the empty set if there is no entry. A history $H(t)$ is the sequence of the actions of all players $< t$. $H(0)$, the history at time 0 is the empty sequence. A strategy of player i , Σ_i is a mapping from histories of length i to actions. We say that a history is finite with length T if after T no players enter. For a finite history of length T we define the payoff of player i , $\pi(H(t))$ for two kinds of games as :

- a) *First Movers Take Ties (FMT):* Given the locations of all players in the history $H(T)$, compute the Jacquard distance between each player and each set on the lattice, allocating $v(S)$ based on the smallest distance and in case of ties allocating $v(S)$ to the player with the smallest index i . The payoff of player i is then the sum of allocated $v(S)$ less a fixed cost F .
- b) *Equal Tie-breaking (ET):* ET is similar to FMT, except in case of ties, ET allocates $v(S)$ equally among all the players with the smallest Jacquard distance to S .

In both games, we add a small technicality to allow for the fact that our games are potentially infinite, especially off the equilibrium path. For histories that are not finite, the payoffs are $-F$ for all players that enter and 0 for the rest.

5. FRONTIER DESCENT STRATEGY

In this section, we describe an easily computable strategy for the FMT game called Frontier Descent. Intuitively, this strategy descends the lattice starting from the top of the lattice until it finds that going any lower would drop revenue below the fixed cost F . This creates a *frontier* of lattice nodes; the algorithm then picks the *best candidate* in the frontier. This algorithm is linear in the size of the lattice, whereas standard backtracking algorithms [1] take time that is exponential in the size of the game.

Unfortunately, we can prove that the Frontier descent algorithm, while faster, does not compute a best response equilibrium strategy. Instead, we prove it provides a *good enough* payoff for all players who enter such that they all break even. Note that this is akin to the concept of satisficing proposed by Simon [12]. From the perspective of algorithmic game theory, this can be considered to be a fast approximation algorithm. While FMT has a fast safe satisficing strategy, we will show that ET has no safe satisficing strategy; even more surprisingly, ET has the property that even with an infinite amount of potential revenue in the lattice there are equilibria in which *no* player enters.

More formally, we use the standard definition of a Subgame Perfect Nash Equilibrium (SPNE) and add the following definitions to capture satisficing in our context:

Definition 1. A history $H(t)$ is blockading if a player at time t has no profitable location to enter even if no player would enter after him. (Note that if $H(t)$ is blockading then $H(t + 1)$ is blockading too).

Definition 2. A safe satisficing strategy (SSS) is a strategy that for any history $H(t)$ that is not blockading finds a location to enter on the lattice for player t such that player t has positive payoff for all continuation histories that have the property that all players $j > t$ earn positive payoffs.

Why do we add the condition that all players j subsequent to i have positive payoffs? We clearly must add some restrictions on subsequent players because otherwise there can be no defense to any move of an earlier player. If player i moves to some node S with descendant revenue $V(S)$, if we have no restrictions an infinite number of players could then perch on S as well. This will make i 's payoff negative. Of course, it will make the payoff for subsequent players (*antagonists*) negative as well. The restriction that the payoff of subsequent players be positive removes these trivial counterexamples and allows reasonable strategies.

Note that the standard definition of an SPNE and Nash Equilibria disallow antagonists because they require that all subsequent players make their best response. However, these are also hard to compute, which is why we are motivated to define a safe satisficing strategy that weakens the standard definition of rationality for subsequent players, but precludes complete irrationality.

With our definition of safe satisficing strategies (SSS) behind us, we now motivate our Frontier Descent algorithm, which is an SSS, by a series of examples. In all examples, assume we are dealing with the First Mover (FMT) game. We will return to the ET game at the end of the section.

The first example shown in Figure 3 motivates the need for descent in the lattice because positioning a player too high in the lat-

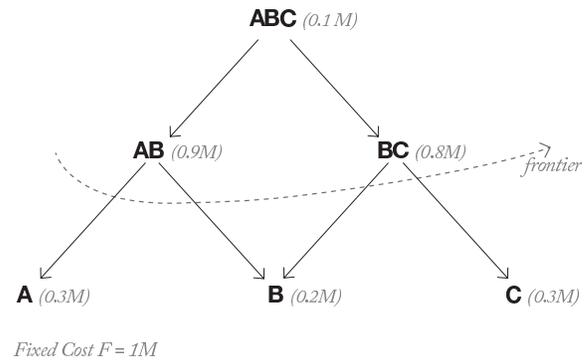


Figure 3: An example that shows that later players can use specialization to undercut the moves of earlier players.

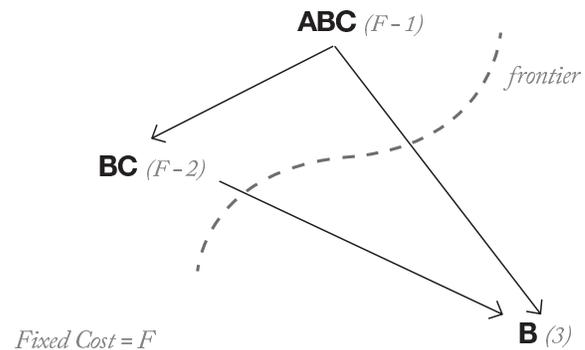


Figure 4: An example that shows how a lower cardinality set can allow a later player to steal from an earlier player

tice can sometimes be a poor long-term strategy. In the figure, suppose Player 1 decides to move to the topmost node ABC . In lieu of other player moves, Player 1 can collect $2.1M$. But in that case, later players, Player 2 and Player 3 can move to nodes AB and BC respectively. Not only is this not an optimal strategy for Player 1, but this is a losing strategy as well! This is because at the end of this history, Player 1 has $0.1M$ which is less than the fixed cost of $1M$.

Reflecting on this example, we see that by picking node ABC , Player 1 exposed itself to more specialized players who leave Player 1 with its original node and no descendants. (The original node is impossible to steal away in the FMT game because no other node can have smaller Jacquard distance).

This suggests that instead of aspiring to optimality, which exists but is most likely hard to compute, Frontier Descent merely tries to protect its descendants against future entrants to ensure a positive payoff. A simple way to do this is to descend the lattice in all directions and keep doing so until one finds a set of *frontier nodes*. Each node S in the frontier must satisfy: 1) $V(S) > F$ and 2) no descendant C of S has $V(C) > F$. Recall that $V(S)$ the descendant revenue also includes the revenues from all descendants of a node while $v(S)$ includes only S 's revenue. The frontier for Figure 3 is shown as a dashed line with the nodes immediately above the dashed line in the set of frontier nodes. The two frontier nodes are thus AB and BC .

Which frontier node should the algorithm pick? It is tempting to try the greedy strategy:

Definition 3. (Greedy Strategy Attempt) The first player picks the largest revenue node in the frontier.

For example, in the Greedy Strategy the first player would pick Node AB in Figure 3. This strategy, however, can fail in other examples as shown in Figure 4. Here, the two frontier nodes are ABC and BC and the greedy strategy should pick ABC for Player 1 because it has higher descendant revenue as $V(ABC) = F - 1 + 3 = F + 2$ compared to $V(BC) = F - 2 + 3 = F + 1$. Unfortunately, Player 2 can pick node BC and simply steal away descendant B ! This is because the Jacquard distance of player 2 from set B , $JD(BC, B) = 1/2$ which is smaller than the Jacquard distance of player 1 from set B which is $JD(ABC, B) = 2/3$.

Thus besides the threat of specialization (attacks from below) one also has to worry about lower cardinality sets (attacks from the side). This suggests a simple modification.

Definition 4. (First player Frontier descent) The first player picks the largest revenue node S among all the smallest cardinality sets in the frontier.

In the analysis, we prove that the first player is immune to all future attacks on its descendant revenue $V(S)$; while the player may get other revenue from other lattice nodes as a bonus, it cannot count on such revenue. However, the descendant revenue $V(S)$ at the time the player located on S is guaranteed at the end of the game.

This guarantee suggests a simple iterative satisficing algorithm. Once a player has picked a node S we simply remove S and all its descendants from the lattice and iterate. This suggests the general algorithm:

General Frontier descent.

Player i picks the largest revenue node S among all the smallest cardinality sets in the frontier of the lattice it starts with (Player 1 starts with the original lattice). S and all its descendants are deleted from the lattice and Player $i + 1$ repeats the algorithm on the reduced lattice. Player 1 starts with the original lattice. The iterations continue until there are no nodes in the frontier of the final reduced lattice.

There are important questions about the efficiency of this algorithm but we can see that even a naive version of this algorithm costs no more than $O(EN)$ where E is the number of edges in the lattice and N is the number of lattice nodes. We will see later that computational efficiency can be improved to $O(E)$ which can be exponentially better than the size of the game tree, which is the standard way to compute an optimal SPNE.

To gain intuition, we examine the General Frontier Descent Algorithm in action on Figure 3. The frontier for the first player is as shown. We pick node AB (highest payoff with $V(AB) = 1.4M$) for Player 1. When we do so, we delete AB , A , and B . This leaves a reduced lattice with ABC , BC and C . The new frontier for the reduced lattice is only BC and this is the node picked for Player 2. The final reduced lattice is only ABC with reduced value $V(ABC) = 0.1$. The final lattice does not possess a frontier and so the algorithm terminates. The final payoff for Player 1 will be $1.5M$ because Player 1 takes the revenue of ABC as first mover. Player 2 gets a revenue of $1.3M$.

We can also quickly show that General Frontier Descent does not always compute an equilibrium as shown in Figure 5. The frontier is as shown. Note that Node ABC is not on the frontier because it has a descendant (Node AB) whose descendant value $V(AB) > F$ by the definition of a frontier. Recall that this was a design decision meant to forestall the threat of specialization! Thus

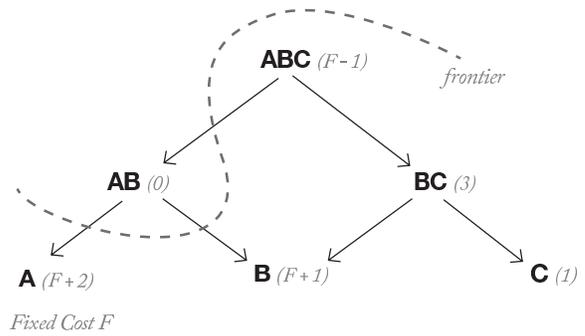


Figure 5: An example that shows that frontier descent does not always compute a SPNE for the FMT game

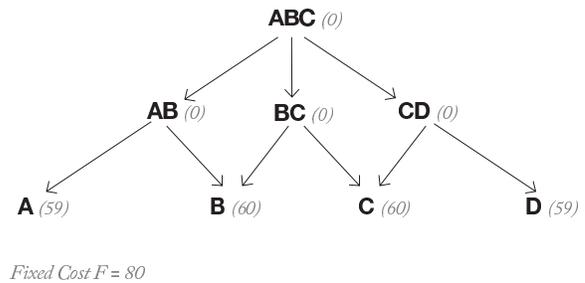


Figure 6: An example that shows that the frontier descent does not produce a satisficing strategy for the ET game.

Player 1 picks Node A . Even in the reduced lattice, ABC is not on the frontier and so Player 2 picks Node B . But in the final reduced lattice, ABC becomes part of the frontier and so Player 3 picks Node ABC . Hence in the final payoffs, Player 3 gets a payoff of 3 while Player 1 gets a payoff of 2.

On the other hand, it is easy to see that there is an SPNE in which Player 1 moves to ABC , Player 2 moves to A , and Player 3 moves to B , leading to payoffs of 3, 2, and 1 respectively. Thus the frontier descent algorithm does not produce an SPNE because Player 1 can improve his lot by playing first at the topmost node. In some sense, Frontier Descent leaves some revenue on the table both from nodes on the top (e.g., ABC) and nodes at the side (e.g., BC and C).

So far we have been talking about the FMT game. It is natural to ask how the Frontier Descent Algorithm does on the ET game. Before even asking the question, we need to modify the FMT frontier descent algorithm. Recall that in the FMT version, after each player located at a node S , the algorithm deletes all descendants of S . This can no longer be done because descendants can now be shared by later players with the same Jacquard distance.

Thus in ET Frontier Descent, we retain all nodes till the end but add a bookkeeping variable to each node with the set of current owners. When a new player i descends the lattice, player i must account for the potential descendant revenue of a node S among all its descendants by sharing equally among all descendants that have the same Jacquard distance. Subject to this modification, a frontier can be calculated for each player, and once again each player can pick the largest revenue node among the smallest cardinality sets in its frontier. Then the ownership sets are updated.

Even with these modifications, ET Frontier Descent does rather badly—in fact, it can lose money as shown in Figure 6. According to frontier descent, nodes AB , BC , and CD are part of the fron-

tier for Player 1. However, all nodes in the frontier have the same cardinality of 2, and BC has the highest value of $V(BC) = 120$. So Frontier Descent picks BC . But picking BC is a bad idea because Player 2 can subsequently locate on AB , and Player 3 can locate on CD . This causes Player 1 a net loss because in the ET game, B is now shared with Player 2, and C is shared with Player 3. Thus Player 1 gets a total revenue of 60 ($30 + 30$) which does not recompense Player 1 for its fixed cost of 80. We will show in the analysis that the ET game has even more unusual properties even in equilibria.

6. ANALYSIS

We start by showing existence of equilibria for both FMT and ET; in particular we show that both games possess not just a Nash Equilibrium but also a subgame perfect Nash equilibrium or SPNE. For readers unfamiliar with the definition of an SPNE, in an SPNE every possible subgame (or subtree in the game tree representing the strategy) of the SPNE is also a Nash Equilibrium.

The reasoning for the existence of an SPNE in both games is similar to the standard proofs for finite extensive form games. A slight difficulty is that our games have an infinite sequence of potential entrants and hence does not have a finite tree. In fact (see the Snowflake example in Figure 7), the threat of an infinite number of entrants makes the equilibrium outcomes quite different from the case when the players know there are a bounded number of entrants.

We surmount this small difficulty by observing that even with an infinite number of *potential* entrants, the number of *actual* entrants must be bounded in any equilibrium and instead of backward induction on the sequence of players we use backward induction over possible entry locations on the lattice.

THEOREM 1. (Equilibrium Existence) *For either of the two games FMT or ET there exists a (generically unique) subgame perfect equilibrium in pure strategies (SPNE).*

PROOF. Consider any history $H(t)$. For any player t the payoff from entering at location S on the lattice is bounded from above by the payoff that player would obtain if there was no more entry after him. If that payoff bound is negative, entry at that location is a dominated strategy (by the action of no entry). For the FMT game any set S which already has one player has the lower bound equal to $-F$. For the ET game, however, even a copy-cat strategy (which locates in a node chosen by an earlier player) can yield a positive payoff. We distinguish two cases.

FMT game:

For any history $H(t)$, let $\gamma(H(t))$ be the number of locations on the lattice that have a strictly positive upper bound. $\gamma(H(t))$ is bounded by $\gamma(\emptyset)$ which in turn is smaller than the size of the lattice. $\gamma(H(t))$ decreases over time (on and off equilibrium path). When $\gamma(H(t)) = 0$ there is a unique continuation SPNE in which no more players enter. Now we use an induction argument. Suppose that for any t and $H(t)$ such that $\gamma(H(t)) \leq M$ there exists a SPNE in pure strategies. For each of these histories select one of these SPNE. Consider any history $H(t')$ such that $\gamma(H(t')) = M + 1$. If a player t' does not enter, he gets a payoff 0. If that player enters at any of the locations with a positive payoff bound, then $\gamma(H(t')) \leq M$ and we have selected a unique SPNE, which allows us to uniquely compute continuation payoffs of player t' upon entry. Entry in any other location is dominated. Since player t' chooses from a finite set of entry locations (plus the option to not enter), there exists an action with a maximum payoff. Pick any of these best response actions as the equilibrium strategy for player t' after history $H(t')$. That implies that for every history $H(t')$ such

that $\gamma(H(t')) = M + 1$ there exists a pure-strategy continuation SPNE. By induction, it is true also for the empty history, and that is the SPNE for the whole game.

ET game:

We need to modify slightly the definition of $\gamma(H(t))$ since a player entering in set S does not exclude the possibility that additional players would enter in that location. Therefore, if given a history $H(t)$ the upper payoff bound for entry in node S is more than F , let the contribution of this node to $\gamma(H(t))$ be equal to the number of players that can enter in that node and make a positive profit assuming no entry anywhere else. That still leaves $\gamma(H(t))$ to be decreasing over time and bounded by $\gamma(\emptyset)$ which in turn is smaller than the number of nodes in the lattice times the ratio $v(L)/F$. The rest of the argument follows without change. \square

We now show that Frontier Descent is a safe satisficing strategy as define earlier. Doing so requires the following lemma that shows that lower cardinality sets can defend against attacks from the side from higher cardinality sets.

LEMMA 1. (Common descendant distance) *If two sets S_1 and S_2 have a common descendant S_3 in the lattice, then: $|S_1| < |S_2|$ if and only if $JD(S_1, S_3) < JD(S_2, S_3)$.*

PROOF. Let x be the cardinality of S_1 , y the cardinality of S_2 , and c the cardinality of common descendant S_3 . We know that S_3 is a subset of S_1 and S_3 is a subset of S_2 . Thus $JD(S_1, S_3) = (x - c)/x = 1 - c/x$. Similarly, $JD(S_2, S_3) = 1 - c/y$. Clearly, if $x < y$, and $x, y, c > 0$, then $c/x > c/y$ and $1 - c/x < 1 - c/y$. The converse holds similarly. \square

THEOREM 2. (FMT safe satisficing) *In the FMT game, Frontier Descent produces a safe satisficing strategy.*

PROOF. We claim that the Frontier Descent Algorithm described above is a safe satisficing strategy. We need to show that when player i makes a move there is a continuation History in which, regardless of the moves of subsequent players $j > i$, the payoff of Player i remains positive. At the time i made its move, $V(S) > F$. We now show that Player i gets a revenue no less than $V(S)$ and a payoff no less than $V(S) - F$. Note that $V(S)$ is the value at the time i made its move. Suppose that some other later player $j > i$ causes Player i 's descendant revenue to drop below $V(S)$. This can only happen if Player j steals a descendant D of S . But that can only happen if D is also a descendant of S' that node j moves to and $JD(S', D) < JD(S, D)$. But in that case (by the Common descendant distance lemma), then S' has smaller cardinality than S . But in that case Player i would have chosen S' instead of S when Player i evaluated its frontier because S would have had payoff greater than F and smaller cardinality than S' . This contradicts the fact that Player i picked the smallest cardinality set in its frontier.

There are two subtleties to this argument. First, the argument tacitly assumes that S' would have been in the frontier when player i evaluated its options. This follows because as the lattice reduces at each stage of the algorithm, nodes like S' can only lower their values of descendant revenue $V(S')$ (monotonicity). Thus if S' was on the frontier at a later stage, it must have been on the frontier at an earlier stage. Second, the argument assumes that Player j cannot pick some descendant R of S . In this case, j will steal R away from i . But by the definition of the frontier, we know that $V(R) < F$. Hence, j will have a negative payoff which implies this is not a satisficing strategy. \square

It is well known [1] that for extended games, one can calculate SPNE using a traversal of the game tree from the bottom up. Since,

we know that FMT has an SPNE and it can be calculated, why bother with a satisficing strategy. This is because the best known general algorithm is linear in the size of the game tree. But the game tree is exponential in the number of nodes N of the lattice L . This is because the first level has N children (Player 1 can move to any node), each child has at least $N - 1$ children (Player 2 can move to any nodes not taken by player 1) and so on. Thus the game tree has $N * (N - 1) * (N - 2) \dots = N!$ nodes which by Stirlings Formula is $O((N/e)^N)$. By contrast, we can show that Frontier Descent can be made to run in $O(E)$, where E is the number of edges in the lattice which is $O(N \log N)$.

THEOREM 3. (Frontier descent is polynomial time) *Frontier descent for FMT can be implemented to run in $O(N \log N)$ time.*

PROOF. We consider a worst-case complete lattice in say n variables and first calculate the number of direct edges E between nodes and immediate descendants. The lattice has $2^n = N$ nodes. Observe that sets of size i have i links to all immediate descendants of size $i - 1$. Thus the total number of links is $E = \sum_{i=0}^n i C(n, i)$ where $C(n, i)$ is the i -th binomial coefficient. By the binomial theorem, $(x + 1)^n = \sum_0^n C(n, i) x^i$. By differentiating both sides and setting $x = 1$ we get the required sum $E = \sum_{i=0}^n i C(n, i) = n 2^{n-1}$. But since $n = \log N$, $E = O(N \log N)$.

The straightforward implementation of FMT Frontier Descent is $O(EN)$. This is because each descent of the lattice can visit edges at most once to compute the $V(S)$ values and the frontier. It can perform at most N descents for at most N players, because there can be at most one player per node.

The running time can be improved by having back pointers from descendants to immediate ancestors and by incrementally recomputing the frontier when a Player i is assigned to a node. We still delete all the player's descendants. However, we also follow the up-pointers to adjust the descendant revenues of all immediate ancestors. If some such ancestors were in the earlier frontier for Player i but are no longer in the frontier for Player $i + 1$, the algorithm has to go further up to *repair* the frontier by following more up-pointers. But once a node is not part of the frontier, it will not re-enter at a later stage of the algorithm.

Thus the frontier repair operation visits every node and every edge at most once. Thus the initial descent is $O(E)$ and the repair is $O(E)$ and so the final algorithm is $O(E)$ which is $N \log N$. \square

We now turn our attention to the ET game. Given that the FMT game may be hard to compute an equilibrium for, it seems reasonable to believe that ET is even harder. Thus it makes sense to look for a satisficing strategy for the ET game. Unfortunately, we can show that:

THEOREM 4. (ET is not safe satisficing) *In the ET game, in general there does not exist a SSS.*

PROOF. For our proof we will exhibit a specific lattice L for which there exists no SSS. More specifically, we will show that for any choice of the location of Player 1 there exists a continuation strategy that is:

- Blockading (recall that a blockading history is one where there is no location at which a player can enter, and still recover its fixed costs, assuming no future moves).
- All players after Player 1 make positive payoffs
- Player 1 makes a negative payoff.

The lattice L consists of 4 locations/nodes: A, B, C and D . No other lattice nodes have any value. There are 60 users in each node. The fixed cost is 80.

There are 4 cases for the first move of Player 1. A location can either be a singleton (sets such as A), a doubleton (sets such as AB), a tripleton (sets such as ABC) or the ground set ($ABCD$).

- Player 1 picks a singleton set. By symmetry, assume Player 1 locates on Node A . Consider the continuation strategy where Player 2 locates on BC and Player 3 locates on CD . Player 1's payoff is -20 (revenue of 60 less a fixed cost of 80) while Players 2 and 3 have a payoff of $90 - 80 = 10$. It is easy to see that this history is blockading. This is because Player 4 cannot locate on any singleton (because Player 4 would at most get 60 from say locating at B); cannot locate on any doubleton (because Player 4 can at most take half of two existing singletons resulting in a revenue of at most 60); cannot locate on either any tripleton or the ground set because the distance to any singleton is larger than that of existing players.
- Player 1 picks a doubleton. By symmetry assume Player 1 enters on Node AB . After the first player enters at AB there is a continuation history in which Players 2 and 3 enter at DA and BC respectively. If they do, they each get a revenue of $60 + 30$, so they make a payoff of 10 each. But Player 1 loses money because it now has revenue of $30 + 30$ and a payoff of -20 . Again, this history is blockading because of a similar case analysis to the one done above.
- Player 1 picks a tripleton set. In this case, Player 2 picks AB and Player 3 picks CD . Player 1 gets a payoff of -80 and Player 2 and 3 get a payoff of 40 each. Again, it is easy that this history is blockading. The analysis for the fourth case (when Player 1 picks $ABCD$) is identical.

Thus for lattice L there is no possible first move of Player 1 that can guarantee positive profits in blockading continuation histories of the ET game. \square

Interestingly, there is an equilibrium strategy in the ET game for lattice L in which Player 1 would enter at AB and only Player 2 would enter in equilibrium in the location CD and consequently Player 1 and 2 would both get a payoff of 40. Note that our counterexample lattice L is similar to the one in Figure 6. However, while there we simply needed a counterexample for Frontier Descent, here we need a counterexample for any algorithm. The symmetry in the counterexample is thus crucial. Note that philosophically it is attacks from the side that one has to worry about from nodes who originally have lower value when the frontier is first evaluated.³

A natural question is how many players will enter each game. The following theorem is immediate for the FMT game:

THEOREM 5. *In the FMT game there exists a satisficing strategy in which at least one player will enter and all revenue will be assigned to some player (i.e, no revenue is left on the table).*

PROOF. We know that if $V(L) > 0$, there exists a frontier for at least the first player in the Frontier Descent algorithm. This is

³Note that the satisficing strategy is harder because it allows more adversarial behaviors. For example, Player 2 can choose to locate at BC in a satisficing strategy after Player 1 locates at AB , thereby accruing a payoff of 10, instead of the optimal strategy of locating at CD , which would give a payoff of 20.

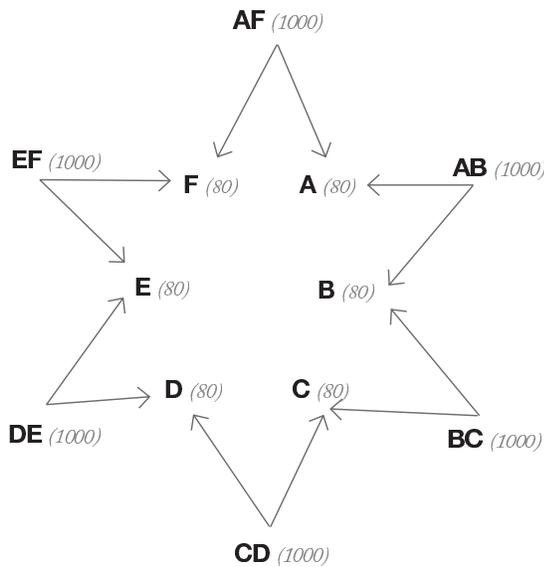


Figure 7: An example that suggests that even large amount of potential revenue there may exist equilibria in the ET game in which no player will enter

because either the top of the lattice is a frontier node or both its descendants are greater than F . In the latter case, we keep descending but maintain the invariant that all ancestors S of visited nodes have $V(S) > F$. But descent must terminate (because we cannot descend beyond the bottom of the lattice). Assume descent terminates at some node E . Then (by the termination condition) none of the immediate descendants D of E have $V(D) > F$. But by the invariant $V(E) > F$. Thus E is a frontier node. Since there must be at least one frontier node, Player 1 will always enter. We have already shown that Frontier Descent is safe satisficing. Finally, note that if at least one player enters, the definition of FMT assigns all revenue to the players who have entered at the end of the game. Thus no revenue is left on the table. \square

Note that it is not possible to provide good bounds on the number of players who enter in the FMT game without making further assumptions on the distribution of total lattice revenue $V(L)$ among individual lattice nodes. For example, consider the lattice where only a single node has value cF for any value of c . No other node has revenue. It is easy to see that only one player enters and takes all the revenue. A slightly less trivial example is where a single node S has value F and S has c immediate descendants of value $F - 1$ and no other lattice nodes have any value. In this case as well, the first player will clean up, accruing revenue $F + c(F - 1)$ for any value of c .

7. OPEN PROBLEMS

We list some interesting open problems suggested by our work.

1. *How many players can enter in an SPNE of the ET game?* The last theorem in the analysis shows that at least one player can enter in the FM game using Frontier Descent. Imagine an ET instance in which the total revenue of the lattice summed across all nodes $V(L)$ is 100 times the fixed cost F . Potentially there is room for 100 entrants. However, we conjecture that there exists ET game instances in which the total potential revenue is arbitrarily high and yet there are SPNEs in which no player enters the game. Further,

if a player deviates and enters, that player could have a negative payoff.

Consider the example in Figure 7. We have shown the lattice using circular rings for each level to show the circularity of the configuration which we refer to as our snowflake example. The lattice on is 6 nodes A through F. All lattice nodes not shown have zero value. The fixed cost is 1100. Assume for a start that Player 1 can only enter at the nodes AB, BC, CD, DE, EF, and AF. By symmetry, assume player 1 enters at AB. In that case, a possible best response for Player 2 is node ED, for player 3 is AF and for player 4 is BC. No more entrants are possible in this history. But in that case AB loses money because his final payoff is 1080 which is less than the fixed cost of 1100. Unfortunately, the complete proof that this is indeed the best response requires arguing over a much larger number of cases. For example, we also need to consider the case when Player 1 enters at node A. We invite the reader to find a proof or a counterexample for our snowflake example!

Note that it is crucial for this example that ET has a possibly infinite set of players who can enter. If we knew that there were only 100 possible players, Player 100, for example, is guaranteed to enter if Players 1 through 99 do not as there can be no subsequent threats from future players. By contrast, FMT has no such problem.

2. *How many players can enter in the FMT game?* We have seen that at least one player can enter in the FMT game and that better bounds are hard to guarantee without further assumptions on how revenue is distributed among nodes. For example, it appears reasonable to assume that interests are distributed according to a power law; this implies that revenues from taste sets also follow power laws. Given such an assumption, it seems possible to prove stronger bounds on the number of entrants.

3. *Are there polynomial time algorithms to compute SPNEs for the FMT or ET games?* We have seen that Frontier Descent takes $O(N \log N)$ but does not guarantee to find the best response. Are there other polynomial time algorithms that either find the best response or find a response that strictly dominates (for the FMT game) Frontier Descent?

4. *Is there a stronger definition of satisficing that allows a satisficing strategy for the ET game?* We have proved that there exists lattices for which the ET game has no satisficing strategy. But our definition of satisficing only requires future players to make money in a continuation history that is blockading. We did not require that the future players also (recursively) play a satisficing strategy. This stronger definition of satisficing—that we call *strong satisficing*—may allow ET to have a satisficing strategy.

5. *What is the relation between strategies that are satisficing and Nash Equilibria?* Consider the following transformed location game where we transform the payoff function π to $\pi' = 1$ if $\pi \geq 0$ and 0 otherwise. In some sense, the new payoff function π' is a first order approximation to the original payoff function. Then it appears that a satisficing solution to the original game is a Nash equilibrium of the transformed game. It also appears that a strong satisficing solution of the original game is an SPNE of the transformed game. These correspondences may make the notion of satisficing less foreign as a notion.

6. *What happens if players could enter in multiple locations?* In the extreme case where each player can enter in an unlimited number of locations, it appears that in any SPNE there would be at most one player earning a strictly positive profit (for generic games). The reasoning is by contradiction: suppose players $t_1, \dots, t_n > 1$ get a positive profit in an equilibrium. Then player 1 could enter in the locations of these players and simply transfer their profits to himself (and at the same time not being threatened by additional entry since we start with an equilibrium).

8. CONCLUSIONS AND FUTURE WORK

This paper has four novel contributions:

- *Modeling Aggregators:* Aggregation is an important and growing phenomenon in the web milieu. We provide the first model of aggregator incentives. The only earlier paper we know of [4] is a paper on the interaction between aggregators and content providers. While this is important, it does not provide insight into why aggregators enter. The two games are orthogonal and should ideally be combined.
- *Jacquard distance for location games:* While we fundamentally model the entrance of aggregators and users as a location game, we are the first to model distance as Jacquard distance and not Euclidean distance. This is not merely an esoteric choice for novelty but, we believe, a more realistic model of user taste sets than Euclidean distance in any dimensional space. We believe our location game, invented to model aggregators, applies to other parts of economics as well in which specialized stores (e.g., grocery stores) steal users from general stores (e.g., Walmart).
- *Safe satisficing strategies:* In the course of this paper, we found it natural to invent safe satisficing strategies. Satisficing is well known in economics; so also are extensive form games. We claim that it is perfectly natural to marry these two concepts. In particular, this is because computing an SPNE is typically exponential in the size of the lattice. To put this into perspective, even for the simple snowflake example with 12 nodes, the size of the lattice is 1024 and the size of the game tree is 1024 factorial (the amount of time required to search the game tree), an astronomically large number. While in practice, there are heuristics such as branch-and-bound, computing SPNEs with the best known algorithms appears very hard. By contrast, we have shown a non-trivial algorithm for the FMT game that is very fast ($N \log N$ versus $N!$) that at least prevents loss if all subsequent players also have positive profits. Could there be simpler game tree strategies for other games as well that are safe satisficing? Recall also our conjecture that a satisficing strategy is a Nash equilibrium of a game with modified payoffs—that are a first order approximation, the sign function, of the original payoff function.

Satisficing strategies can also be viewed from the lens of computer science. Just as there exist polynomial time approximation algorithms for so-called NP complete algorithms (whose best known solutions are exponential in the input), perhaps satisficing algorithms introduce another notion of approximation to algorithmic game theory, in addition to the standard ϵ -Nash.
- *FMT and ET games:* We have made progress on the analysis of aggregator games. We have shown that the FMT game has an $N \log N$ satisficing solution which is not an SPNE. At least one player will enter if the revenue is more than the fixed cost, and no revenue is wasted. We have shown by contrast that the ET game has no satisficing solution.

In conclusion, we note that Frontier Descent provides some intuition into the evolution of aggregation on the real web. In some sense, the sequence of entries from say Yahoo to TechCrunch to AppleInsider is a form of descent. And yet large aggregators such as Yahoo remain profitable (as do general stores like Walmart) perhaps because they aggregate the long tail of content that does not

warrant a sufficiently large market for a more specialized aggregator.

Frontier descent also suggests that as the fixed cost F goes down, more aggregators will enter in more specific niches because players can descend lower in the lattice and still make a profit. Arguably, fixed costs have reduced over the last 5 years because of cloud services and better abstractions for building web sites. This may explain the recent emergence of a large number of aggregators catering to more specific tastes. More importantly, it suggests an interesting business opportunity. A vendor that can provide good tools to reduce the cost F of doing business is likely to open the floodgates for new small aggregators to cater to the long tail of user interests—and reap a rich reward in doing so.

9. REFERENCES

- [1] In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 21. Cambridge University Press, 2007.
- [2] A. Ansari, N. Economides, and J. Steckel. The max-min-min principle of product differentiation. *Journal of Regional Science*, 38(2):207–230, 1998.
- [3] A. Caplin and B. Nalebuff. Aggregation and imperfect competition: On the existence of equilibrium. *Econometrica*, 59(1):25–59, 1991.
- [4] C. Dellarocas, Z. Katona, and W. Rand. Media, aggregators and the link economy: Strategic hyperlink formation in content networks. *Boston University School of Management Research Paper, Number 2010-30*, 2010.
- [5] N. Economides. Symmetric equilibrium existence and optimality in differentiated product markets. *Journal of Economic Theory*, 47(1):178–194, 1989.
- [6] H. Hotelling. Stability in competition. *The Economic Journal*, 39(153):41–57, 1929.
- [7] R. P. McAfee, J. McMillan, and M. D. Whinston. Multiproduct monopoly, commodity bundling, and correlation of values. *The Quarterly Journal of Economics*, 104(2):371–383, 1989.
- [8] B. Nalebuff. Bundling as an entry barrier. *The Quarterly Journal of Economics*, 119(1):159–187, 2004.
- [9] S. C. Salop. Monopolistic competition with outside goods. *Bell Journal of Economics*, 10:141–156, 1979.
- [10] R. Schmalensee. Gaussian demand and commodity bundling. *The Journal of Business*, 57(1):S211–S230, 1984.
- [11] A. Shaked and J. Sutton. Relaxing price competition through product differentiation. *The Review of Economic Studies*, 49(1):3–13, 1982.
- [12] H. A. Simon. Rational choice and the structure of the environment. *Psychological Review*, 63(2):129–138, 1956.
- [13] G. J. Stigler. United states v. loew’s inc.: A note on block-booking. *The Supreme Court Review*, 1963:152–157, 1963.