Dynamic Inefficiency: Anarchy without Stability

Noam Berger* Michal Feldman† Ofer Neiman‡ Mishael Rosenthal§

May 8, 2011

Abstract

The price of anarchy [17] is by now a standard measure for quantifying the inefficiency introduced in games due to selfish behavior, and is defined as the ratio between the optimal outcome and the worst Nash equilibrium. However, this notion is well defined only for games that always possess a Nash equilibrium (NE). We propose the *dynamic inefficiency* measure, which is roughly defined as the average inefficiency in an infinite best-response dynamic. Both the price of anarchy [17] and the price of sinking [9] can be obtained as special cases of the dynamic inefficiency measure. We consider three natural best-response dynamic rules — *Random Walk* (RW), *Round Robin* (RR) and *Best Improvement* (BI) — which are distinguished according to the order in which players apply best-response moves.

In order to make the above concrete, we use the proposed measure to study the job scheduling setting introduced in [3], and in particular the scheduling policy introduced there. While the proposed policy achieves the best possible price of anarchy with respect to a pure NE, the game induced by the proposed policy may admit no pure NE, thus the *dynamic inefficiency* measure reflects the worst case inefficiency better. We show that the dynamic inefficiency may be arbitrarily higher than the price of anarchy, in any of the three dynamic rules. As the dynamic inefficiency of the RW dynamic coincides with the *price of sinking*, this result resolves an open question raised in [3].

We further use the proposed measure to study the inefficiency of the Hotelling game and the facility location game. We find that using different dynamic rules may yield diverse inefficiency outcomes; moreover, it seems that no single dynamic rule is superior to another.

1 Introduction

Best-response dynamics are central in the theory of games. The celebrated Nash equilibrium solution concept is implicitly based on the assumption that players follow best-response dynamics until they reach a state from which no player can improve her utility. Best-response dynamics give rise to many interesting questions which have been extensively studied in the literature. Most of the focus concerning best response dynamics has been devoted to convergence issues, such as whether best-response dynamics converge to a Nash equilibrium and what is the rate of convergence.

Best-response dynamics are essentially a large family of dynamics, which differ from each other in the order in which turns are assigned to players [1]. It is well known that the order of the players' moves is crucial to various aspects, such as convergence rate to a Nash equilibrium [6]. Our main goal is to study the effect of the players' order on the obtained (in)efficiency of the outcome.

The most established measure of inefficiency of games is the Price of Anarchy (PoA) [14, 17], which is a worst-case measure, defined as the ratio between the worst Nash equilibrium (NE) and the social optimum (with respect to a well-defined social objective function), usually defined with respect to pure strategies. The PoA essentially measures how much the society suffers from players who maximize their individual welfare rather than the social good. The PoA has been evaluated in many settings, such as selfish routing [21, 20], job scheduling [14, 5, 8], network formation [7, 1, 2], facility location [22], and more. However, this notion is well defined only in settings that are guaranteed to admit a NE.

*Einstein Institute of Mathematics, Hebrew University of Jerusalem. Email: berger@math.huji.ac.il
†School of Business Administration and Center for Rationality, Hebrew University of Jerusalem. Email: mfeldman@huji.ac.il
‡Princeton university and Center for Computational Intractability. Email: oneiman@princeton.edu
§School of Engineering and Computer Science, Hebrew University of Jerusalem. Email: mishael@cs.huji.ac.il
[1] In fact, best-response dynamics may be asynchronous, but in this paper we restrict attention to synchronous dynamics.

One approach that has been taken with respect to this challenge, in cases where agents are assumed to use pure strategies, is the introduction of a *sink equilibrium* [9]. A sink equilibrium is a strongly connected component with no outgoing edges in the *configuration graph* associated with a game. The configuration graph has a vertex set associated with the set of pure strategy profiles, and its edges correspond to best-response moves. Unlike pure strategy Nash equilibria, sink equilibria are guaranteed to exist. The social value associated with a sink equilibrium is the expected social value of the stationary distribution of a random walk on the states of the sink. The *price of sinking* is the equivalence of the price of anarchy measure with respect to sink equilibria.

Indeed, the notion of best response lies at the heart of many of the proposed solution concepts, even if just implicitly. The implicit assumption that underlies the notion of social value associated with a sink equilibrium is that in each turn a player is chosen uniformly at random to perform her best response. However, there could be other natural best-response dynamics that arise in different settings.

In this paper, we focus on the following three natural dynamic rules: (i) **random walk (RW)**, where a player is chosen uniformly at random; (ii) **round robin (RR)**, where players play in a cyclic manner according to a pre-defined order, and (iii) **best improvement (BI)**, where the player with the current highest (multiplicative) improvement factor plays. Our goal is to study the effect of the players' order on the obtained (in)efficiency of the outcome.

To this end, we introduce the concept of *dynamic inefficiency* as an equivalent measure to the price of anarchy in games that may not admit a Nash equilibrium or in games in which best-response dynamics are not guaranteed to converge to a Nash equilibrium. Every dynamic rule D chooses (deterministically or randomly) a player that performs her best response in each time period. Given a dynamic D and an initial configuration u, one can compute the expected social value obtained by following the rules of dynamic D starting from u. The dynamic inefficiency in a particular game is defined as the expected social welfare with respect to the worst initial configuration. Similarly, the dynamic inefficiency of a particular family of games is defined as the worst dynamic inefficiency over all games in the family. Note that the above definition coincides with the original price of anarchy measure for games in which best-response dynamics always converge to a Nash equilibrium (e.g., in congestion [19] and potential [15] games) for every dynamic rule. Similarly, the dynamic inefficiency of the RW dynamic coincides with the definition of the price of sinking. Thus, we find the dynamic inefficiency a natural generalization of well-established inefficiency measures.

1.1 Our Results

We evaluate the dynamic inefficiency with respect to the three dynamic rules specified above, and in three different applications, namely non-preemptive job scheduling on unrelated machines [3], the Hotelling model [11], and facility location [16]. Our contribution is conceptual as well as technical. First, we introduce a measure which allows us to evaluate the inefficiency of a particular dynamic even if it does not lead to a Nash equilibrium. Second, we develop proof techniques for providing lower bounds for the three dynamic rules. In what follows we present our results in the specific models.

Job scheduling (Section 3) We consider job scheduling on unrelated machines, where each of the n players controls a single job and selects a machine among a set of m machines. The assignment of job i on machine j is associated with a processing time that is denoted by $p_{i,j}$. Each machine schedules its jobs sequentially according to some non-preemptive scheduling policy (i.e., jobs are processed without interference, and no delay is introduced between two consecutive jobs), and the cost of each job in a given profile is its completion time on its machine. The social cost of a given profile is the maximal completion time of any job (known as the makespan objective).

Machines' ordering policies may be local or strongly local. A local policy considers only the parameters of the jobs assigned to it, while a strongly local policy considers only the processing time of the jobs assigned to it on itself (without knowing the processing time of its jobs on other machines). Azar et. al. [3] showed that the PoA of any local policy is $\Omega(\log m)$ and that the PoA of any strongly local policy is $\Omega(m)$ (if a Nash equilibrium exists). Ibarra and Kim [12] showed that the shortest-first (strongly local) policy exhibits a matching $O(m)$ bound, and Azar et. al. [3] showed that the inefficiency-based (local) policy (defined in Section 3) exhibits a matching $O(\log m)$ bound.

We claim that there is a fundamental difference between the last two results. The shortest-first policy induces a potential game [13, 15]; thus best-response dynamics always converge to a pure Nash equilibrium, and the PoA is an appropriate measure. In contrast, the inefficiency-based policy induces a game that does not necessarily admit a pure

Nash equilibrium [3], and even if a Nash equilibrium exists, not every best-response dynamic converges to a Nash equilibrium. Consequently, the realized inefficiency of the last policy may be much higher than the bound provided by the price of anarchy measure.

We study the dynamic inefficiency of the inefficiency based policy with respect to our three dynamic rules. We show a lower bound of $\Omega(\log\log \mathbf{n})$ for the dynamic inefficiency of the RW rule. This bound may be arbitrarily higher[2] than the price of anarchy, which is bounded by $O(\log \mathbf{m})$. This resolves an open question raised in [3]. For the BI and RR rules, we show even higher lower bounds of $\Omega(\sqrt{n})$ and $\Omega(n)$, respectively.

Hotelling model (Appendix A) Hotelling [11] devised a model where customers are distributed evenly along a line, and there are m strategic players, each choosing a location on the line, with the objective of maximizing the number of customers whose location is closer to her than to any other player. It is well known that this game admits a pure Nash equilibrium if and only if the number of players is different than three. This motivates the evaluation of the dynamic inefficiency measure in settings with three players. The social objective function we consider here is the minimal utility over all players, i.e., we wish to maximize the minimal number of customers one attracts.

We show that the dynamic inefficiency of the BI rule is upper bounded by a universal constant, while the dynamic inefficiency of the RW and RR rules is lower bounded by $\Omega(n)$, where n is the number of possible locations of players. Thus, the BI dynamics and the RW and RR dynamics exhibit the best possible and worst possible inefficiencies, respectively (up to a constant factor). In contrast to the BI dynamics, the RW and RR dynamics are *configuration-oblivious* (i.e., the next move is determined independently of the current configuration).

Facility location (Appendix B) In facility location games a central designer decides where to locate a public facility on a line, and each player has a single point representing her ideal location. Suppose that the cost associated with each player is the squared distance of her ideal location to the actual placement of the facility, and that we wish to minimize the average cost of the players. Under this objective function the optimal location is the mean of all the points. However, for any chosen location, there will be a player who can decrease her distance from the chosen location by reporting a false ideal location. Moreover, it is easy to see that if the players know in advance that the mean point of the reported locations is chosen, then every player who is given a turn can actually transfer the location to be exactly at her ideal point. Thus, unless all players are located at exactly the same point, there will be no Nash equilibrium. Our results indicate that the dynamic inefficiency of the RR and RW rules is exactly 2, while that of the BI rule is $\Theta(n)$.

2 Preliminaries

In our analysis it will be convenient to use the following graph-theoretic notation: we think of the configuration set (i.e., pure strategy profiles) as the vertex set of a *configuration graph* $G = (V, E)$. The configuration graph is a directed graph in which there is a directed edge $e = (u, v) \in E$ if and only if there is a player whose best response given the configuration u leads to the configuration v. We assume that each player has a unique best response for each vertex. A *sink* is a vertex with no outgoing edges. A *Nash Equilibrium* (NE) is a configuration in which each player is best responding to the actions of the other players. Thus, a NE is a sink in the configuration graph.

A *social value* of $S(v)$ is associated with each vertex $v \in V$. Two examples of social value functions are the social welfare function, defined as the sum of the players' utilities, and the max-min function, defined as the minimum utility of any player.

A best-response dynamic rule is a function $D : V \times \mathbb{N} \to [n]$, mapping each point in time, possibly depending on the current configuration, to a player $i \in [n]$ who is the next player to apply her best-response strategy. The function D may be deterministic or non-deterministic.

Let $P = \langle u_1, \ldots, u_T \rangle$ denote a finite path in the configuration graph (where u_i may equal u_j for some $i \neq j$). The *average social value* associated with a path P is defined as $S(P) = \frac{1}{T}\sum_{t=1}^{T} S(u_t)$. Given a tuple $\langle u, D \rangle$ of a vertex $u \in V$ and a dynamic rule D, let $\mathcal{P}_T(u, D)$ denote the distribution over the paths of length T initiated at vertex

[2] Note the parameter n; i.e., number of players, versus the parameter m; i.e., number of machines.

u under the dynamic rule D. The social value of the dynamic rule D initiated at vertex u is defined as

$$S(u, D) = \lim_{T\to\infty} \mathbb{E}_{P\sim\mathcal{P}_T(u,D)}[S(P)]. \quad (1)$$

While the expression above is not always well defined, in Appendix C we demonstrate that it is always well defined for the dynamic rules considered in this paper.

With this, we are ready to define the notion of *dynamic inefficiency*. Given a finite configuration graph $G = (V, E)$ and a dynamic rule D, the dynamic inefficiency (DI) of G with respect to D is defined as

$$\mathrm{DI}(D, G) = \max_{u\in V} \frac{\mathrm{OPT}}{S(u, D)},$$

where $\mathrm{OPT} = \max_{u\in V} S(u)$. That is, DI measures the ratio between the optimal outcome and the social value obtained by a dynamic rule D under the worst possible initial vertex. Finally, for a family of games $\mathcal{G}$, we define the dynamic inefficiency of a dynamic rule D as the dynamic inefficiency of the worst possible $G \in \mathcal{G}$. This is given by

$$\mathrm{DI}(D) = \sup_{G\in\mathcal{G}} \{\mathrm{DI}(D, G)\} .$$

In some of the settings, social costs are considered rather than social value. In these cases, the necessary obvious adjustments should be made. In particular, $S(u, D)$ will denote the social *cost*, OPT will be defined as $\min_{u\in V} S(u)$, and the dynamic inefficiency of some dynamic D will be defined as $\mathrm{DI}(D) = \max_{u\in V} \frac{S(u,D)}{\mathrm{OPT}}$. We consider both cases in the sequel.

An important observation is that both the price of anarchy and the price of sinking are obtained as special cases of the dynamic inefficiency. In games for which every best-response dynamic converges to a Nash equilibrium (e.g., potential games [15]), the dynamic inefficiency is independent of the dynamic and is equivalent to the price of anarchy. The price of sinking [9] is equivalent to the dynamic inefficiency with respect to the RW dynamic rule.

3 Dynamic Inefficiency in Job Scheduling

Consider a non-preemptive job scheduling setting on unrelated machines, as described in the Introduction. Define the efficiency of a job i on machine j as

$$\mathrm{eff}(i, j) = \frac{p_{ij}}{\min_{k\in[m]} p_{ik}}.$$

The efficiency-based policy of a machine (proposed by [3]) orders its jobs according to their efficiency, from low to high efficiency, where ties are broken arbitrarily in a pre-defined way.

A configuration of a job scheduling game is a mapping $u : [n] \to [m]$ that maps each job to a machine. The processing time of machine j in configuration u is $\mathrm{time}_u(j) = \sum_{i\in u^{-1}(j)} p_{ij}$, and the social value function we are interested in is the *makespan* — the longest processing time on any machine, i.e., $S(u) = \max_{j\in[m]} \mathrm{time}_u(j)$.

The players are the jobs to be processed, their actions are the machines they choose to run on, and the cost of a job is its own completion time. Let $u(i)$ be the completion time of job i in configuration u, and let $v \in V$ be the resulting configuration of a best-response move of job i in configuration u. Define $\mathrm{gain}_u(i) = \frac{u(i)}{v(i)}$ (note that the gain is always at least 1).

3.1 Random Walk Dynamic

In this section we consider the RW dynamic, where in each turn a player is chosen uniformly at random to play. The main result of this section is the establishment of a lower bound of $\Omega(\log \log n)$ for the dynamic inefficiency of job scheduling under the efficiency-based policy. This means that the inefficiency may tend to infinity with the number of jobs, even though the number of machines is constant. This result should be contrasted with the $O(\log m)$ upper bound on the price of anarchy, established by [3]. The main result is cast in the following theorem.

Theorem 1. *There exists a family of instances G_n of machine scheduling on a constant number of machines, such that*

$$\mathrm{DI}(\textsf{\textit{RW}}, G_n) \geq \Omega(\log \log n) \ ,$$

where n is the number of jobs. In particular the dynamic inefficiency is not bounded with respect to the number of machines.

Remark: The definition of the dynamic inefficiency with respect to the $\textsf{RW}$ dynamic rule coincides with the definition of the *price of sinking*. Thus, the last result can be interpreted as a lower bound on the price of sinking.

The assertion of Theorem 1 is established in the following sections.

3.1.1 The Construction

Let us begin with an informal description of the example. As the base for our construction we use an instance, given in [3], with a constant number of machines and jobs, that admits no Nash equilibrium. Then we add to it n jobs indexed by $1, \ldots, n$ and one machine, such that the n additional jobs have an incentive to run on two machines. On the first machine, denoted by W, the total processing time of all the n jobs is smaller than 2, while on the second machine T the processing time of any job is ≈ 1. Each of these jobs has an incentive to move from machine W to machine T if it has the minimal index on T, thus increasing the processing time on T. We show that the expected number of jobs on T, and hence also the expected makespan, is at least $\Omega(\log \log n)$, while the optimum is some universal constant.

Formally, there will be 5 machines, denoted by A, B, C, T, W, and $n + 5$ jobs denoted by $0, 1, 2, \ldots, n$ and $\alpha, \beta, \gamma, \delta$. The following table shows the processing time for the jobs on the machines:

	A	B	C	T	W
0	4	24	3.95	25	∞
α	2	12	1.98	∞	∞
β	5	28	4.9	∞	∞
γ	20	∞	∞	∞	∞
δ	∞	∞	50	∞	∞
i	∞	∞	$\frac{1}{50i^3}$	$\sum_{j=1}^{i} \frac{1}{j^2} - \epsilon$	$\frac{1}{i^2}$

In the last row, i stands for any job $1, 2, \ldots, n$, and $\epsilon = \frac{1}{10 \cdot 2^n}$.

3.1.2 Useful Properties

Observation 1. *The inefficiency policy induces the following order on the machines:*

- *On machine A the order is $(\gamma, \alpha, 0, \beta)$.*
- *On machine B the order is $(\beta, \alpha, 0)$.*
- *On machine C the order is $(\delta, \alpha, \beta, 0, 1, 2, \ldots, n)$.*
- *On machine T the order is $(0, 1, 2, \ldots, n)$.*
- *On machine W the order is $(1, 2, \ldots, n)$.*

Proof. On machine C every job has efficiency 1; hence given any tie-breaking rule between jobs of equal efficiency we let δ be the one that runs first[3] (the rest of the ordering is arbitrary). On the other machines, the order follows from straightforward calculations. □

[3]We can also handle tie-breaking rules that consider the length of the job. For instance, if shorter jobs were scheduled first in case of a tie, we would split job δ into many small jobs.

Note that no job except δ has an incentive to move to machine C, and job γ will always be in machine A. The possible configurations for jobs $0, \alpha, \beta$ are denoted by XYZ; for instance, ABB means that job 0 is on machine A and jobs α, β are on machine B. We shall only consider configurations in G that have incoming edges, and in this example there are 8 such configurations among the possible configurations for jobs $0, \alpha, \beta$. The transitions are

$$BAA \to BBA \to ABA \to ABB \to AAB \to TAB \to TAA$$

From state TAA we can go either to TBA or back to BAA. From TBA the only possible transition will take us back to the state ABA. These are all the possible transitions; hence there is no stable state. At any time a job i for $i > 0$ has an incentive to be in machine T only if it has the *minimal* index from all the jobs in T. We want to show that in the single (non-trivial) strongly connected component of G, the expected number of jobs in machine T is at least $\Omega(\log \log n)$.

3.1.3 Dynamic Inefficiency - Lower Bound

It can be checked that any configuration for jobs $1, 2, \ldots, n$ is possible among machines T, W. Consider the stationary distribution π over this strongly connected component in G. Let $T(i) \subseteq V$ be the set of configurations in which job i is scheduled on machine T, and let $p_n(i) = \sum_{v \in T(i)} \pi(v)$ denote the probability that job i is on machine T. Let $p_n(\emptyset)$ be the probability that no job is on machine T.

Observation 2. *For any $n > m \geq i$, $p_n(i) = p_m(i)$. This is because the incentives for jobs $\alpha, \beta, \gamma, \delta, 0, 1, \ldots, i$ are not affected by the presence of any job j for $j > i$.*

Using this observation we shall omit the subscript and write only $p(i)$. The following claim suggests we should focus our attention on how often machine T is empty.

Claim 3. *For any $n \geq 1$, $p_n(\emptyset) \leq p(n)$.*

Proof. Job n has an incentive to move to T if and *only if* the configuration is such that T is empty. The probability of job n to get a turn to move is $1/(n+5)$. Hence the probability that job n is in T at some time t is equal to the probability that for some $i \geq 0$ job n entered T at time $t - i$ (i.e., machine T was empty and n got a turn to play) and stayed there for i rounds. The probability that job n stayed in machine T for i rounds is at least $p_i = \left(1 - \frac{1}{n+5}\right)^i$. We conclude that

$$p(n) \geq \frac{p_n(\emptyset)}{n+5} \sum_{i=0}^{\infty} p_i = p_n(\emptyset).$$

□

The main technical lemma is the following:

Lemma 4. *There exists a universal constant c such that for any $n > 1$, $p_n(\emptyset) \geq \frac{c}{n \log n}$.*

Let us first show that given this lemma we can easily prove the main theorem:

Proof of Theorem 1. In the single (non-trivial) strongly connected component of G, the expected number of jobs in machine T (with respect to a RW) is at least

$$\sum_{i=1}^{n} p(i) \geq c \sum_{i=2}^{n} \frac{1}{i \log i} \geq (c/2) \log \log n\ .$$

However, there is a configuration in which every job completes execution in time at most 50;[4] hence $\mathrm{OPT}(G) \leq 50$. We conclude that the dynamic inefficiency for G is at least $\Omega(\log \log n)$. □

[4]For instance, if all jobs $1, 2, \ldots, n$ are on machine W, then job i will finish in time $\sum_{j=1}^{i} \frac{1}{j^2} < 2$

In what follows, we establish the assertion of Lemma 4.

Proof of Lemma 4. Let $Y \subseteq V$ be the set of configurations in which T is empty, and let t be the expected number of steps between visits to configurations in Y. We have that $p_n(\emptyset) = \frac{1}{t}$, and need to prove that $t \leq O(n \log n)$. We start with a claim on rapidly decreasing integer random variables.

Claim 5. *Fix some $n \in \mathbb{N}$, $n > 1$. Let $x_1, x_2, \ldots$ be random variables getting non-increasing values in $\mathbb{N}$, such that $\mathbb{E}[x_1] \leq n/2$ and for every $i > 0$, $\mathbb{E}[x_{i+1} \mid x_i = k] \leq k/2$; then if we let s be a random variable which is the minimal index such that $x_s = 0$, then $\mathbb{E}[s] \leq \log n + 2$.*

Proof. First we prove by induction on i that $\mathbb{E}[x_i] \leq \frac{n}{2^i}$. This holds for $i = 1$; assume it is true for i and then prove for $i + 1$. By the rule of conditional probability,

$$\begin{aligned} \mathbb{E}[x_{i+1}] &= \sum_{j \geq 0} \Pr[x_i = j] \cdot \mathbb{E}[x_{i+1} \mid x_i = j] \\ &\leq \sum_{j \geq 0} \Pr[x_i = j] \cdot (j/2) = \mathbb{E}[x_i]/2 \leq \frac{n}{2^{i+1}} \,. \end{aligned}$$

Note that if $x_i = 0$ for some i, then it must be that $x_j = 0$ for all $j > i$. Now for any integer $i > 0$,

$$\begin{aligned} \Pr[s = \log n + i] &\leq \Pr[s > \log n + i - 1] \\ &= \Pr[x_{\log n + i - 1} \geq 1] \\ &\leq \mathbb{E}[x_{\log n + i - 1}] \leq 1/2^{i-1} \end{aligned}$$

the second inequality is a Markov inequality. We conclude that

$$\begin{aligned} \mathbb{E}[s] &= \sum_{i=1}^{\log n} i \cdot \Pr[s = i] + \sum_{i=\log n + 1}^{\infty} i \cdot \Pr[s = i] \\ &\leq \log n + \sum_{i=\log n + 1}^{\infty} i/2^{i-1} \leq \log n + 2 \,. \end{aligned}$$

□

Claim 6. *Assume that we are in a configuration u in which job $i \in \{0, 1, \ldots, n\}$ is in machine T. The expected time until we reach a configuration in which job i is not in machine T is at most $O(n)$.*

Proof. First note that $p(0) = c'$ for some constant $0 < c' < 1$, this is because jobs $0, \alpha, \beta, \gamma, \delta$ are not affected at all by the location of any job from $1, 2, \ldots, n$, and hence when they get to play they will follow one of the two cycles shown earlier, which implies that in a constant fraction c' of the time job 0 will be in machine T, and in the other $1 - c'$ fraction it will be on another machine.

Note that job i will have incentive to leave T if job 0 is in machine T when i gets its turn to play. Let $q(i)$ be the event that i gets a turn to play (which is independent of the current configuration), $T(i)$ denotes the event that job i is in machine T, then we have that the probability that job i will leave machine T is at least

$$\begin{aligned} \Pr[q(i) \wedge T(0) \mid T(i)] &= \Pr[q(i) \mid T(0) \wedge T(i)] \cdot \Pr[T(0) \mid T(i)] \\ &= Pr[q(i)] \cdot \Pr[T(0)] = \frac{c'}{n+5} \end{aligned}$$

Now the expected time until job i will leave is at most $\frac{1}{\Pr[q(i) \wedge T(0) | T(i)]} \leq O(n)$. □

Claim 7. *Let $\ell = \ell(t)$ be a random variable that is the minimal job in T at time t, and let x be the random variable that is the next job that enters T. Then $\mathbb{E}[x \mid \ell = m] \leq m/2$.*

Proof. The jobs that have an incentive to enter machine T are $0, 1, \ldots, m-1$. It is easy to see that for any job $i \in \{1, \ldots, m-1\}$, $\Pr[x = i \mid \ell = m] \leq \frac{1}{m-1}$ (note that job 0 has also some small probability to enter T, but it does not contribute to the expectation). Now

$$\mathbb{E}[x \mid \ell = m] = \sum_{i=1}^{m-1} i \cdot \Pr[x = i \mid \ell = m] \leq m/2 \; .$$

□

Let $y \in Y$ be any configuration in which T is empty. We define a series of random variables $x_1, x_2, \ldots$ as follows. Let x_1 be the index of the first job to enter T, and let x_i be the maximal index of a job in T when x_{i-1} leaves T (and 0 if T is empty). Note that when $x_{i-1} = k$ is the maximal index of a job in T, no job with index larger than k has an incentive to move to T; hence given that $x_{i-1} = k$ it must be that $x_i \leq k$. Let s be the minimal index such that $x_s = 0$; then either machine T is empty or by Claim 6 is expected to become empty in $O(n)$ steps. It is easy to see that $\mathbb{E}[x_1] \leq n/2$, and the tricky part is to bound the expected value of x_i.

Claim 8. $\mathbb{E}[x_i | x_{i-1} = k] \leq k/2$.

Proof. Fix some k such that $x_{i-1} = k$. Consider the time $\bar{t}$ in which job k became the maximal job in T, and consider the time $t' < \bar{t}$ in which job k moved to machine T and did not leave until time $\bar{t}$. In time t' it must be that no job i, for $i < k$, was in machine T, since job k had an incentive to move to T. In particular, x_i is not in T at time t'.

Consider the time $t'' > t'$ in which x_i enters T, and stays until job k leaves. In time t'' the minimal job in T is at most k; hence by Claim 7 we have that $\mathbb{E}[x_i \mid x_{i-1} = k] \leq k/2$. □

Consider the random variables $x_1, x_2, \ldots, x_s$. By Claim 8 they satisfy the conditions of Claim 5; hence $\mathbb{E}[s] \leq \log n + 2$. By Claim 6 we have that the expected time until the maximal job leaves T is at most $O(n)$. We conclude that in expectation after $O(n \log n)$ steps the maximal job in T will be 0, and within additional $O(n)$ steps machine T is expected to become empty. This concludes the proof.

□

3.2 Best Improvement Dynamic

In this section we show a lower bound on the dynamic inefficiency with respect to the BI dynamic; specifically we show for an infinite number of $n \in \mathbb{N}$ an instance G_n of job scheduling on n jobs and five machines such that

$$\mathrm{DI}(\mathsf{BI}, G_n) \geq \Omega(\sqrt{n}) \; .$$

3.2.1 The Construction

	A	B	C	T	Q
0	4	24	3.95	25	∞
α	2	12	1.98	∞	∞
β	5	28	4.9	∞	∞
γ	20	∞	∞	∞	∞
δ	∞	∞	50	∞	∞
i	∞	∞	$\frac{1}{50\sqrt{n}\cdot 4^{i-1}}$	$\frac{1}{\sqrt{n}}$	$\frac{1}{e^{i-1}}$

For $1 \leq k \leq n$ denote by $U_k \subseteq V$ the set of configurations in which jobs $1, \ldots, k$ are on machine Q, jobs $k+1, \ldots, n$ are on machine T, and job 0 is not on machine T. Denote by $V_k \subseteq V$ the set of configurations in which jobs $1, \ldots, k$ are on machine T, jobs $k+1, \ldots, n$ are on machine Q, and job 0 is not on machine T.

Informally, the dynamics of this game are such that starting in some configuration in U_n, all jobs $1, \ldots, n$ will move to machine T; i.e., move to configuration in V_n; then the jobs will start to move back to Q, starting with job n, then $n-1$, etc. However there is a number $1 \leq x \leq n$ such that any configuration in V_x is stable for jobs $1, \ldots, n$

(none of them has an incentive to move). Then, similarly to the previous example, jobs $0, \alpha, \beta, \gamma, \delta$ will start playing, and after at most 7 steps job 0 reaches machine T; at this point all the rest of the jobs $1, \dots, x$ will have the higher gain and will move to machine Q. At this point jobs $0, \alpha, \beta, \gamma, \delta$ will play again and after at most 3 steps job 0 leaves machine T, which brings us back to a configuration in U_n.

The idea is that in this process there are at most $2n + 10$ steps, and in a constant fraction of them the finish time of machine T is at least $\sqrt{n}/2$, which means that the dynamic inefficiency is at least $\Omega(\sqrt{n})$ (it is easy to see that $OPT \leq 50$).

3.2.2 Dynamic Inefficiency - Lower Bound

To show this formally, we begin with some observations:

1. For any $0 \leq i < j \leq n$, $\text{eff}(i, T) < \text{eff}(j, T)$ and for any $1 \leq i < j \leq n$, $\text{eff}(i, Q) < \text{eff}(j, Q)$.
2. If job 0 is on machine T, then the next steps will move all the other jobs on machine T to Q.

Proof of observation 2. It can be checked that for jobs $0, \alpha, \beta, \gamma, \delta$ the maximal gain is smaller than 2, and as the maximal finish time on Q is $e/(e-1) < 2$, when job 0 is on machine T any other job on T has gain larger than $\frac{25}{2} > 2$ to move to Q. (Note that any job in Q has no incentive to move to T since its finish time will be larger than 25 if it does so). □

We proceed with several claims that together establish the desired lower bound.

Claim 9. *For any $k \geq 1$ and configuration $u \in U_k$, $\text{gain}_u(k) > \text{gain}_u(z)$ for any job $z \neq k$.*

Proof. Let $k \geq 1$ be an integer and $u \in U_k$. We have that $\text{gain}_u(k) > \text{gain}_u(i) \geq \sqrt{n}$ for any $0 < i < k$, since any job $i \in \{1, \dots, k\}$ will finish in time $1/\sqrt{n}$ if it chooses to move to T and has finish time at least 1 on machine Q. Also job k has the highest finish time on machine Q, hence it has the highest gain among $\{1, \dots, k\}$.

It remains to show that all other jobs have gain smaller than $\sqrt{n}$. The finish time of any job $j > k$ on machine T is strictly smaller than $\sqrt{n}$, and if moves to Q, the first observation suggests that it will finish after job 1 (whose run time on Q is 1), so its gain is strictly smaller than $\sqrt{n}$. All the other jobs $0, \alpha, \beta, \gamma, \delta$ have at most constant gain. □

Let $x = W(e\sqrt{n})$ be the solution to the following equation

$$xe^{x-1} = \sqrt{n} ,$$

where $W(x)$ is the W-Lambart function. W.l.o.g. assume that $x \in \mathbb{N}$.[5]

Claim 10. *For any $x < k \leq n$ and configuration $u \in V_k$, job k will play next.*

Proof. It is sufficient to show that $\text{gain}_u(k) > \text{gain}_u(z)$ for any job $z \neq k$. Since $k > x$ we have that $k \cdot e^{k-1} > \sqrt{n}$, hence $\text{gain}_u(k) = \frac{k/\sqrt{n}}{1/e^{k-1}} > e$. First we show that for $1 \leq i < k$, $\text{gain}_u(k) > \text{gain}_u(i)$, this is because $\text{gain}_u(k) = k \cdot e^{k-1}/\sqrt{n} > i \cdot e^{i-1}/\sqrt{n} = \text{gain}_u(i)$. For $i \geq k+1$ we have that

$$\text{gain}_u(i) = \frac{\sum_{j=i}^{n} 1/e^{j-1}}{k/\sqrt{n}} < \frac{\sqrt{n}}{k \cdot e^{i-2}} \leq \frac{\sqrt{n}}{k \cdot e^{k-1}} \leq 1 .$$

As for jobs $0, \alpha, \beta, \gamma, \delta$ we noted that their maximal gain is at most 2. □

Claim 11. *Configuration V_x is stable for jobs $1, \dots, n$.*

Proof. From jobs $1, \dots, x$ job x has the highest finish time in T and the lowest finish time if it moves to Q, and job x has no incentive to move to Q since its finish time in T is $x/\sqrt{n} = 1/e^{x-1}$ which is its finish time on Q. Among jobs $x+1, \dots, n$ job n would have the highest incentive, since its finish time in Q is the highest (any of them would finish in time $(x+1)/\sqrt{n}$ moving to T), but job n finish time on Q is $\sum_{j=x+1}^{n} 1/e^{j-1} < 1/e^{x-1} = x/\sqrt{n}$, hence it has no incentive to move and finish in time larger than $x/\sqrt{n}$. □

[5] More careful calculation can show that it is enough to assume that $x - \lfloor x \rfloor < 1/2$.

We conclude that when starting in any configuration $u \in U_n$, Claim 9 implies that the next n steps will have jobs $n, n-1, \ldots, 1$ moving to machine T. Claim 10 and the definition of V_k implies that the next $n-x$ steps will move jobs $n, n-1, \ldots, x+1$ back to machine Q. By Claim 11 in configuration V_x jobs $1, \ldots, n$ will not move, thus only jobs $0, \alpha, \beta, \gamma, \delta$ will play, and once job 0 arrive to machine T (after at most 7 steps) all jobs other than 0 will leave T back to Q (by observation 2). At this point job 0 will leave T in at most 3 steps and we are back to a configuration in U_n.

3.3 Round Robin Dynamic

In this section we consider a dynamic in which there is a pre-defined ordering on the players, and they take turns accordingly. We will show an example in which the dynamic inefficiency is $\Omega(n)$.

3.3.1 The Construction

The example is similar to the example in the previous section. The ordering on the jobs is $n, n-1, \ldots, 1, \delta, \gamma, \beta, \alpha, 0$ and the initial configuration is: all jobs $1, \ldots, n$ on machine Q, jobs $0, \beta$ on machine A, and job α on B (as always job γ will be on A and δ on C).

	A	B	C	T	Q
0	4	24	3.95	25	∞
α	2	12	1.98	∞	∞
β	5	28	4.9	∞	∞
γ	20	∞	∞	∞	∞
δ	∞	∞	50	∞	∞
i	∞	∞	$\frac{1}{50n \cdot 4^{i-1}}$	$1/2$	$\frac{1}{e^{i-1}}$

Recall that the transitions for jobs $0, \alpha, \beta$ starting at ABA are such that when they are given a turn according to the RR rule:

$$ABA \to ABB \to AAB \to TAB,$$

and then when they are given another turn

$$TAB \to TAA \to TBA \to ABA\,.$$

Now after one round all jobs $1, \ldots, n$ will move to T, starting from job n and concluding with job 1, and then job 0 will move to T as well. In the next round all jobs $1, \ldots, n$ will have incentive to leave T back to Q, and then job 0 will leave T as well. Thus we return to the initial configuration. Note that in those $2n + 10$ turns, at least n of them have $n/2$ jobs on machine T, hence the average finish time in the RR dynamic is at least $n/10$, while OPT has a constant makespan.

4 Conclusion

We study the notion of dynamic inefficiency, which generalizes well-studied notions of inefficiency such as the price of anarchy and the price of sinking, and quantify it in three different applications. In games where best-response dynamics are not guaranteed to converge to an equilibrium, dynamic inefficiency reflects better the inefficiency that may arise. It would be of interest to quantify the dynamic inefficiency in additional applications. It is of a particular interest to study whether there exist families of games for which one dynamic rule is always superior to another.

In the job scheduling realm, our work demonstrates that the inefficiency based policy suggested by [3] suffers from an extremely high price of sinking. A natural open question arises: is there a local policy that always admits a Nash equilibrium and exhibits a PoA of $o(m)$? (recall that m is the number of machines). Alternatively, is there a local policy that exhibits a dynamic inefficiency of $o(m)$ for some best response dynamic rule? Recently, Caragiannis [4] found a *preemptive* local policy that always admits a Nash equilibrium and has a PoA of $O(\log \mathbf{m})$. However, we are primarily interested in non-preemptive policies.

References

[1] S. Albers, S. Elits, E. Even-Dar, Y. Mansour, and L. Roditty. On Nash equilibria for a network creation game. In *seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2006.

[2] E. Anshelevich, A. Dasgupta, É. Tardos, and T. Wexler. Near-Optimal Network Design with Selfish Agents. In *STOC'03*, 2003.

[3] Y. Azar, K. Jain, and V. Mirrokni. (almost) optimal coordination mechanisms for unrelated machine scheduling. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '08, pages 323–332, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.

[4] I. Caragiannis. Efficient coordination mechanisms for unrelated machine scheduling. In *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '09, pages 815–824, Philadelphia, PA, USA, 2009. Society for Industrial and Applied Mathematics.

[5] A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. In *SODA*, pages 413–420, 2002.

[6] E. Even-Dar, A. Kesselman, and Y. Mansour. Convergence time to nash equilibria. In *ICALP*, pages 502–513, 2003.

[7] A. Fabrikant, A. Luthra, E. Maneva, C. Papadimitriou, and S. Shenker. On a network creation game. In *ACM Symposium on Principles of Distributed Computing (PODC)*, 2003.

[8] M. Feldman and T. Tamir. Conflicting congestion effects in resource allocation games. In *Proceedings of the 4th International Workshop on Internet and Network Economics*, WINE '08, pages 109–117, Berlin, Heidelberg, 2008. Springer-Verlag.

[9] M. Goemans, V. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *FOCS '05: Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 142–154, Washington, DC, USA, 2005. IEEE Computer Society.

[10] R. Holzman. An axiomatic approach to location on networks. *Mathematics of Operations Research*, 15(3):553–563, 1990.

[11] H. Hotelling. Stability in competition. *Economic Journal*, 39(53):41–57, 1929.

[12] O. H. Ibarra and C. E. Kim. Heuristic algorithms for scheduling independent tasks on nonidentical processors. *Journal of the ACM*, 24:280–289, 1977.

[13] N. Immorlica, L. Li, V. S. Mirrokni, and A. S. Schulz. Coordination mechanisms for selfish scheduling. *Theor. Comput. Sci.*, 410:1589–1598, April 2009.

[14] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *STACS*, pages 404–413, 1999.

[15] D. Monderer and L. S. Shapley. Potential Games. *Games and Economic Behavior*, 14:124–143, 1996.

[16] Herve Moulin. On strategy-proofness and single-peakedness. *Public Choice*, 35:437–455, 1980.

[17] C. Papadimitriou. Algorithms, games, and the Internet. In *Proceedings of 33rd STOC*, pages 749–753, 2001.

[18] Eric Rasmusen. *Games and information: an introduction to game theory*. Blackwell Publishers, 2001.

[19] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[20] T. Roughgarden. The price of anarchy is independent of the network topology. In *STOC'02*, pages 428–437, 2002.

[21] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236 – 259, 2002.

[22] A. R. Vetta. Nash equilibria in competitive societies with applications to facility location, traffic routing and auctions. In *Symposium on the Foundations of Computer Science (FOCS)*, pages 416–425, 2002.

A Hotelling Model

In this section we study the dynamic inefficiency in a family of games induces by the Hotelling model described in the introduction. We show that the choice of the dynamic rule greatly affects the dynamic inefficiency. Specifically, we show that when the dynamic D is one of RR or RW then $\mathrm{DI}(D) = \Theta(n)$ (where n is the number of possible locations for the players), which is the worst asymptotically possible, while if D is the BI dynamic then $\mathrm{DI}(D) = \Theta(1)$, namely the best asymptotically possible.

Assume that the customers are evenly distributed on the interval $[0, 1]$, and let the set $L = \{i/(n-1)\}_{i=0}^{n-1}$ denote the possible player locations. Each vertex in the configuration graph $G = (V, E)$ is an injective mapping of $[m]$ to L. In what follows we refer to the location $i/(n-1) \in L$ as *location* i. For a given vertex $u \in V$, the vector $u = (u_1, \dots, u_m)$ denotes the respective locations of the m players. Let u^s be the sorted (in ascending order) vector u, and let i^s be the index of player i in u^s. Then, the utility of a player $i \in [m]$ in the strategy profile associated with u is $\frac{u^s_{i^s+1} - u^s_{i^s-1}}{2}$ (where we fix $u^s_0 = 0$ and $u^s_{m+1} = 1$). We assume that if a player has more than a single best response move, she will break ties by choosing the location with the largest index among equal-utility locations[6].

It is well known that an Hotelling setting with three players admits no pure Nash equilibrium (see, e.g., in [18], chapter 14). This is also a setting in which best response dynamics are very natural. Understanding the effect of various dynamic rules under this setting is, therefore, very interesting. Therefore from now on we consider the case $m = 3$. We begin by presenting a general claim that holds for every dynamic rule D without starvation[7].

Claim 12. *There exists a time $I = I(D) \in \mathbb{N}$ such that for every time step $t \geq I$, there are two players in consecutive locations, i.e. there exists an $0 \leq i < n-1$ such that two players are in locations $i, i+1$.* [8]

Proof. Since there is no NE in three players game and since there is no starvation, eventually one player will get a turn and change her location. Assume that the two players that do not play next are in locations j, k, then it is easy to verify that the possible best-response locations of the third player are one of $j-1, j+1, k-1, k+1$. □

In the following sections we denote by I the time step after which there are two consecutive players in every period.

A.1 Random Walk Dynamic

In this section we show that giving turns to players uniformly at random usually results in configurations in which one of the players gains very little. Let $\mathrm{span}_u(t)$ be a random variable whose value is the distance between the leftmost and the rightmost players after t steps starting in configuration u (and multiplied by $n-1$, for normalization).

Claim 13. *For every step $t > I$, the span of the players is 2 with probability at least $1/3$.*

Proof. By Claim 12 in any configuration at time $t \geq I$ there exists an i such that two players are at locations $i, i+1$. With probability $1/3$ the turn will be given to the third player, and it is easy to check that her best response is moving either to location $i-1$ or to $i+2$. □

Claim 14. *In any configuration, every move can increase the span of the players by at most 1.*

[6] Our results will still hold for any deterministic pre-defined tie breaking rule.

[7] A dynamic rule is said to have no starvation if for any $T \in \mathbb{N}$ and for every player i it holds that $\Pr(\text{player i receives a turn in some time } t > T) = 1$.

[8] For random dynamics this may hold only with high probability.

Proof. It is easy to see that the leftmost or rightmost players will have no incentive to increase their distance, and that if the middle player has incentive to move it will be one location to the right of the rightmost player, or one location to the left of the leftmost player, increasing the span by 1. $\square$

Lemma 15. *For any initial configuration u there exists $T \in \mathbb{N}$ such that for any $t > T$, $\mathbb{E}[\text{span}_u(t)] \leq 5$.*

Proof. For any $t > I$ by Claim 13 and Claim 14 we have

$$\mathbb{E}[\text{span}_u(t)] \leq (1/3) \cdot 2 + (2/3) \cdot (\mathbb{E}[\text{span}_u(t-1)] + 1)$$

For any $c > 0$, if it is the case that in time $t > I$ we have that $\mathbb{E}[\text{span}_u(t)] \leq 4+c$, then $\mathbb{E}[\text{span}_u(t+1)] \leq 4+2c/3$. We conclude that the expected span will decrease to below 5 after at most $T = I + \log_{2/3} n$ turns, and remain smaller than 5 forever after. $\square$

Note that the span of the players is an upper bound on the minimal gain (the middle player gains half of the span divided by $n-1$), so Lemma 15 implies that the dynamic inefficiency for the RW dynamic is $\Omega(n)$: In OPT the minimal gain is $1/3$, while the expected minimal gain in the RW dynamic is bounded by $5/(2(n-1))$.

A.2 Round Robin Dynamic

In this section we show that giving turns to players by the RR rule may once again result in a poor performance with respect to OPT. Assume w.l.o.g that n is odd, and that the three players $1, 2, 3$ play in this order. Let $i = (n-1)/2$ and let $u = (i/(n-1), (i-1)/(n-1), (i+1)/(n-1))$ be the initial configuration (*i.e.* player 1 in location i, player 2 in location $i-1$ and 3 in location $i+1$). The next 12 turns will go over the following configurations.

$$\begin{aligned}
&(i, i-1, i+1) \to (\mathbf{i{+}2}, i-1, i+1) \to (i+2, \mathbf{i}, i+1) \\
&\to (i+2, i, \mathbf{i{-}1}) \to (\mathbf{i{+}1}, i, i-1) \to (i+1, \mathbf{i{+}2}, i-1) \\
&\to (i+1, i+2, \mathbf{i}) \to (\mathbf{i{-}1}, i+2, i) \to (i-1, \mathbf{i{+}1}, i) \\
&\to (i-1, i+1, \mathbf{i{+}2}) \to (\mathbf{i}, i+1, i+2) \to (i, \mathbf{i{-}1}, i+2) \\
&\to (i, i-1, \mathbf{i{+}1}),
\end{aligned}$$

thus we return to the initial configuration, and hence will stay in this loop forever. In every configuration of this loop the minimal gain is smaller than $2/(n-1)$, and as the social value at OPT is $1/3$ the dynamic inefficiency is $\Omega(n)$.

A.3 Best-Improvement Dynamic

In this section we show that the BI dynamic does well with respect to OPT, in contrast to the previous dynamics we studied. In fact the average minimal gain among the three players is shown to be at least $1/10$, where $\text{OPT} = 1/3$.

Let $A = 1 - \sqrt{1/2}$ and $B = \sqrt{1/2}$, assume w.l.o.g that $A \approx a/(n-1)$ and $B \approx b/(n-1)$ for some integers a, b, so we refer to A, B as valid locations for players. The proof that BI dynamic does well will include two steps. First we show that at some point all three players will have locations inside the interval $[A, B]$. The next step shows that the players will always stay inside $[A, B]$, while keep increasing the minimal gain, and then at some point the best improvement move will put the three players in adjacent locations either near A or near B.

For a configuration v denote by $b(v)$ the minimal distance of any player from the boundary, *i.e.* if $v = (j, k, \ell)$ then $b(v) = \min\{j, k, \ell, 1-j, 1-k, 1-\ell\}$.

Claim 16. *For any configuration v, if $b(v) < A$ then after several turns we will arrive to a configuration w such that $b(w) > b(v)$.*

Proof. By Claim 12 we may assume that we eventually reach to a configuration v such that the players are in $j, k-1, k$ with $j < k$. Assume w.l.o.g that $b(v) = j$ and that $j < 1/2$. Consider the following options:

1. If the player at k has the highest incentive to play, it can be checked that her move will be to location $k-2$.

2. Otherwise, if $k < \frac{1+\sqrt{1+8j}}{4}$ then it can be checked that the player at $k-1$ will move to location $k+1$. The assumption that $j < A$ implies that $\frac{1+\sqrt{1+8j}}{4} \leq B$, hence $k+1 \leq B$.

3. Otherwise, if $k \geq \frac{1+\sqrt{1+8j}}{4}$ then the player at j will have the highest incentive to move, and she will go to $k-2$, which will reduce $b(\cdot)$ and we are done[9].

$\square$

Claim 16 immediately implies that after several turns we will arrive to a configuration v in which $b(v) \geq A$, *i.e.* all players will be inside the interval $[A, B]$. Now to analyze the minimal gain starting from such a configuration[10]. Intuitively the minimal gain will keep increasing by $1/(2(n-1))$ at every turn, until we reach a configuration with two players at locations A, B. At this point the best improvement move will reduce the minimal gain back to $1/(n-1)$.

More formally, assume we are in such a configuration $u = (j, k, \ell)$ (w.l.o.g $A \leq j < k < \ell \leq B$). It is easy to check that the minimal gain is obtained by the player at location k, and unless $j = A$ and $\ell = B$ the middle player also has the highest incentive to play[11], her best response move is either $j-1, \ell+1$, thus increasing by $1/(2(n-1))$ the minimal gain (note that she will remain in the interval because one of locations A, B is not taken). Once we reach a configuration $(A, B-1, B)$ (this is w.l.o.g by symmetry), where the minimal gain is $\frac{B-A}{2}$, the highest incentive will be of the player at A to move to $B-2$, then the minimal gain goes down to $1/(n-1)$. Note that it will take $(n-1)\cdot(B-A)-2$ turns in order to increase the minimal gain back to $\frac{B-A}{2}$, in each step the increase is by $1/(2(n-1))$, so in average over this sequence of turns the minimal gain is $\frac{B-A}{4} = \frac{\sqrt{2}-1}{4} \approx 0.1$.

Remark: In the BI dynamic rule the choice of the next player is based (amongst other factors) on the current configuration. This is in contrast to the RR and RW rules, where the choice of the next player is independent of the current configuration. This may lead one to wonder whether the dependence on the current configuration is a necessary condition for avoiding a high dynamic inefficiency. In the following section we show a deterministic dynamic rule which is oblivious to the current configuration (and also to the initial configuration) and yet achieves a dynamic inefficiency of a constant.

A.4 A Good Configuration-Oblivious dynamic rule

Denote the players by $1, 2, 3$, then the dynamic rule is $((123)^n(123)^n(132)^m(123)^n)^*$, where $m = n$ if $n \mod 4 \neq 0$ and $m = n-1$ otherwise. Denote by R the sequence $(123)^n(123)^n(132)^m(123)^n$. There are $\approx 12n$ turns in every cycle, and the cycles are played infinitely.

Lemma 17. *In every cycle there are at least $\Omega(n)$ turns in which the minimal gain is at least $\Omega(1)$.*

It is easy to see that proving this lemma will give the result we desire - that the average minimal gain is at least $\Omega(1)$. Assume w.l.o.g that $4|n-1$, and let $a = 1/4$ and $b = 3/4$.

Sketch of Lemma 17. The first step of the proof shows that after $4n$ turns (playing R once) all of the players will be located inside the interval $[a, b]$. The proof is very similar to that of Claim 16.

Informally, there are three possible behaviors once the players are inside $[a, b]$, they can either spread out, tend to the center, or remain at the center.

- While remaining in the center all three players will be at the interval $[1/2 - 2/(n-1), 1/2 + 2/(n-1)]$ and while playing in the same order (like 123) they will not leave the interval, and repeat the same configuration every 4 cycles of 123 (or 132), this is exactly what happens in the bad example for RR described in Section A.2.

[9] Note that it could be that $j = 1-k$ in which case the value does not drop, however then it can be checked that the next case to happen will be Case 2., also note that Case 2. cannot be followed by Case 1., and when Case 3. will finally happen, the value of $b(\cdot)$ will surely drop (since $k \leq B$).

[10] In the calculation of the dynamic inefficiency, the value of any finite number of first rounds does not matter

[11] A and B were chosen such that if indeed $j = A$, $\ell = B$, n is large enough and w.l.o.g (by Claim 12) $k = \ell - 1$, then the gain of the player at j moving to $k-1$ is slightly larger than $\frac{B}{(B+A)/2}$, while the gain of the player at k moving to $\ell+1$ is slightly smaller than $\frac{A}{(B-A)/2}$, which are equal.

- While spreading out the locations will be either: $(j, j+1, k)$, where $j = 1-k$, and the player at $j+1$ will have a turn and move to location $k+1$. (Or equivalently $(j, k-1, k)$ with $j = 1-k$, and the player at location $k-1$ will move to location $j-1$). The rotation will be such that the middle player always gets a turn to play and moves just outside of the "opposite" location. The spreading stops when one of two things happen: either the ordering has changed, or two players are located on the "boundaries" at a, b, then the middle player has no incentive to move and the isolated player will "close the gap" and move to location $a+2$ (or $b-2$). Once spreading is stopped, the tending to the center will occur.

- While tending to the center, the locations will be $(j-1, j, j+1)$, and the player at $j-1$ will move to $j+2$ (or player at $j+1$ move to $j-2$, depending on whether j is smaller or bigger than $1/2$). This behavior will stop when $j \in [1/2 - 1/(n-1), 1/2 + 1/(n-1)]$, then we can either spread again or remain at the center.

The purpose of the first $(123)^n$ steps, is to bring the players to the center, they will surely arrive to the center at some point during this stage, because after spreading (which takes $n/2$ turns at most) the next phase is tending to the center. The purpose of the next $(123)^n$ steps is to continue and spread out the players, however this may not happen (it will if the tending to the center happened from the left side ($j < 1/2$). If the tending happened from the right then the next $(132)^m$ steps may spread the players. There are four possible configurations in which the order may change once the players have reached the center, and they will spread when we change the order in three of those. The purpose of the last $(123)^n$ is to spread the players if we happened to be in that fourth configuration when the order was changed (note that since $m \mod 4 \neq 0$ we will arrive to a different configuration when going back to the order (123), so it can be checked that a spread is guaranteed).

We conclude that in every cycle of R there will be at least one spread out phase, in which the minimal gain keep increasing by $1/(2(n-1))$ for $\Omega(n)$ time, hence in some constant fraction of the cycle the minimal gain is $\Omega(1)$. □

B Facility Location

In a facility location setting, a central designer wishes to locate a public facility at some point on the line $\mathbb{R}$. Each player $i \in [n]$ has a true preferred location x_i^t, referred as her *ideal location*. Each player declares some ideal location x_i, and the designer chooses a location $y = \mu = f(x) = f(x_1, \ldots, x_n)$, according to some function f. In the simplest version, the cost of each player $i \in [n]$ is given by the distance $|x_i^t - y|$, and the designer's objective is to minimize the sum of the players' costs. It is well known [16] that choosing the median of the n locations minimizes the objective function and is also incentive compatible, i.e., it is always in the best interest of every player i to declare her true ideal location, independent of the other players' declarations.

We consider a different objective function — that of minimizing the sum of squared distances, i.e., $\sum_{i=1}^{n}(x_i - y)^2$ (See Holzman [10] for an axiomatic characterization of this objective function). It is not too difficult to verify that the solution to this minimization problem is the mean of $\{x_1, \ldots, x_n\}$, namely $y = \frac{\sum_{i=1}^{n} x_i}{n}$. The resulting sum of squared distances is then $n \cdot var(x) = \sum_{i=1}^{n}(x_i - \mu)^2$. Unlike the last example, a mechanism choosing the mean point is not incentive compatible. Furthermore, for any declaration x_{-i} of all other players except i, player i can ensure that the chosen location will be x_i^t by falsely declaring $x_i = nx_i^t - \sum_{j \neq i} x_j$.

In the facility location model the strategy space of each player is infinite, thus Claim 21 cannot be directly applied. Nevertheless, since the chosen location will always be one of $\{x_1^t, \ldots, x_n^t\}$, $S(u, D)$ and $DI(D)$ are well defined for the dynamic rules above.

We next show that $\mathrm{DI}(\mathsf{RR}) = \mathrm{DI}(\mathsf{RW}) = 2$ while $\mathrm{DI}(\mathsf{BI}) = \Theta(n)$.

Lemma 18. *The dynamic inefficiency of the* RR *and* RW *dynamic rules in facility location settings is 2.*

Proof. w.l.o.g we assume $\mu^t = \frac{\sum_{i=1}^{n} x_i^t}{n} = 0$, otherwise shift all the locations by a constant. Hence $OPT = n \cdot var(x^t) = \sum_{i=1}^{n}(x_i^t - \mu^t)^2 = \sum_{i=1}^{n}(x_i^t)^2$. On the other hand in each of the above dynamics $y = x_i^t$ an expected $\frac{1}{n}$

of the time. It follows that the expected sum of cost is:

$$\begin{aligned}\frac{1}{n}\sum_{i,j}(x_i^t - x_j^t)^2 &= 2\sum_{i=1}^{n}(x_i^t)^2 - \frac{2}{n}\sum_{i,j}x_i^t \cdot x_j^t \\ &= 2\sum_{i=1}^{n}(x_i^t)^2 - \frac{2}{n}\sum_{i}x_i^t\sum_{j}x_j^t \\ &= 2\sum_{i=1}^{n}(x_i^t)^2 = 2\cdot OPT,\end{aligned}$$

where the third equality follows from the fact that $\sum_j x_j^t = 0$. □

We next present an upper bound for the dynamic inefficiency of facility location games with respect to any dynamic rule.

Lemma 19. *The dynamic inefficiency of any dynamic rule D is bounded by $O(n)$.*

Proof. Assume w.l.o.g. that $x_1^t \le x_2^t \le \ldots \le x_n^t$ and that $\mu^t = 0$. Assume also that $|x_1^t| \ge |x_n^t|$. Since for player i playing her best response insures that $y = x_i^t$, it follows that the expected sum of cost cannot exceed:

$$\begin{aligned}\sum_i (x_1^t - x_i^t)^2 &= \sum_i (x_1^t)^2 + (x_i^t)^2 - 2x_1^t \cdot x_i^t \\ &= n(x_1^t)^2 + OPT - 2x_1^t \cdot \sum_i x_i^t \\ &\le n(x_1^t)^2 + OPT \le (n+1)OPT,\end{aligned}$$

where the last inequality follows from the fact that $OPT \ge (x_1^t)^2$. □

Remark: The cost of player i may be zero, in which case the (multiplicative) best-improvement is not well defined. Naturally, we would like to assume that if the cost of player i is larger than the cost of player j when y is the chosen location, then player i will gain more than player j by reducing her cost to zero[12]. We conclude that the BI dynamic rule starves all the players except the extreme two indefinitely.

A linear dynamic inefficiency follows.

Lemma 20. *The dynamic inefficiency of the BI dynamic rule in facility location settings is $\Theta(n)$.*

Proof. We examine the following example: $x_1^t = -1$, $x_2^t, \ldots, x_{n-1}^t = 0$, $x_n^t = 1$. Hence $OPT = 2$. On the other hand under the BI dynamic $y \in \{x_1^t, x_n^t\}$ at all times. Thus the expected sum of costs is $n - 2 + 4 = n + 2$. □

C DI is Well Defined

Claim 21. *Consider a finite configuration graph G. The social value $S(u, D)$ on G is always well defined for $D \in \{RW, RR, BI\}$.*

Proof. Consider first the deterministic RR rule. Since the number of configurations is finite, and at each configuration the turn can be given to one of n players, at some point in time, we must revisit some vertex $v_1 \in V$ with the turn given to the same player $i \in [n]$. Let $(v_1, \ldots, v_k, v_1)$ denote the cycle starting at vertex v_1 until revisiting v_1. Since the RR dynamic is deterministic, this cycle will repeat indefinitely, and the obtained social value is

$$S(u, D) = \frac{1}{k}\sum_{j=1}^{k} S(v_j)\ .$$

[12]This can be formulated by adding a small constant ϵ to the costs of all players.

The argument for the BI dynamic rule is even simpler, as the player to play next is determined by the configuration alone.

It remains to show that the social value is well defined with respect to the RW rule. Let $\mathcal{S}$ denote the set of all sink equilibria of the configuration graph. The RW dynamic rule induces a Markov chain on the vertices of the configuration graph. Since once some vertex of a given sink is visited, the random walk never leaves that sink, we can define for each tuple $\langle T, u \rangle$ of a sink equilibrium $T \in \mathcal{S}$ and an initial vertex $u \in V$, the probability $P_u(T)$ that a random walk starting at u will enter the sink T. Also, since every walk will eventually reach some sink equilibrium with probability 1, it holds that $\sum_{T \in \mathcal{S}} P_u(T) = 1$. Since every sink equilibrium is strongly connected and every vertex in a sink equilibrium has a self loop (the best response of the player who best responded in the last turn is the same vertex), the induced chain on T is ergodic, hence admitting a stationary distribution π_T. The social value is therefore given by

$$S(u, D) = \sum_{T \in \mathcal{S}} P_u(T) \left(\sum_{v \in T} \pi_T(v) S(v) \right).$$

$\square$