The Price of Fairness of Scheduling a Scarce Resource

Shai Vardi
Purdue University, West Lafayette, Indiana 47907, svardi@purdue.edu

William B. Haskell
Purdue University, West Lafayette, Indiana 47907, whaskell@purdue.edu

Scarce resources are often shared between different stakeholders, and need to be scheduled fairly among them. In this paper, we study the utility loss resulting from requiring the schedule to be fair, using the *price of fairness* – the ratio of the utility attainable with and without fairness restrictions. We focus on envy-freeness, an intuitive and well-studied fairness notion. We derive tight bounds on the price of fairness as a function of the problem parameters – the number of agents, the time horizon, the discount factor, the switching cost (the time it takes to transfer possession of the resource), and the concavity and heterogeneity of the agents' utility functions.

We analyze the effect of the different parameters on the utility loss to obtain actionable strategic insights. At the macro level, the price of fairness is increasing in the number of agents, the switching cost, and the heterogeneity of the utility functions, and it is decreasing in the time horizon and the degree of concavity. At the micro level, however, the price of fairness is not monotone in any of these parameters, except the heterogeneity. This implies that, counter-intuitively, a scheduler may be able to reduce the loss of utility by *increasing* the number of agents or *decreasing* the time horizon. Furthermore, we find that the dependence of the loss of utility on the number of agents obeys a threshold rule. These insights can help guide decision-makers in scheduling problems where fairness is a factor.

1. Introduction

In many operational settings, a central planner needs to schedule a resource among different entities, who use it to create value. The goal of the central planner is maximize the total value generated; however different circumstances may dictate that the schedule meet some fairness requirement. If the resource was procured jointly, the contributing entities would wish to guarantee that they each capture their fair share of the welfare; an organization may wish to generate fair schedules to enhance employee satisfaction and retention (Bianchi et al. 2015); a government may wish to schedule a scarce resource in a fair fashion among its states or districts. To make this problem more concrete, consider the following examples.

- *Scientific instruments.* Some scientific instruments, such as particle colliders, quantum computers, and telescopes, are extremely expensive and are constructed/procured jointly by multiple parties. For example, the Large Binocular Telescope (LBT) is a joint international project between multiple institutions. The demand for observation far exceeds the supply (Johnston 1990), and while the goal is to maximize the scientific benefit,[1] the joint ownership constrains the schedule planner to guarantee sufficient observation time to all stakeholders.
- *Corporate car sharing.* Having a smaller fleet of cars that are shared among employees is an effective way for companies to reduce their fleet-related costs. It may be suboptimal to schedule the fleet to maximize the total utility derived from the fleet, as a schedule that is perceived as unfair would lead to employee dissatisfaction.
- *Healthcare.* Advanced medical equipment is usually scarce and needs to be scheduled among different medical teams. For example, several surgeons often share a surgery room which they jointly rent and for which they jointly procured the equipment. Because of the shared ownership, the room needs to be scheduled fairly among the surgeons, even if one of them has a higher success rate.
- *Project scheduling.* Skilled workers/machines are in short supply but need to be allocated among multiple simultaneous projects. For example, a consultant has been jointly contracted for several projects but can only work on one at a time; specialty construction equipment is needed for several projects, but it is costly to transfer it between projects.

In all of the above situations, the fairness requirement comes at a cost, as it imposes additional constraints, which can adversely affect the total achievable welfare. Several notions of fairness are studied in the literature, including envy-freeness, proportionality, equitability, and max-min fairness (see the discussion of related literature below). We focus on *envy-freeness*, a well studied

[1] The scientific benefit is determined through a peer-evaluation process, see e.g., Shetrone et al. (2007).

fairness notion, which has been called the "gold standard" of fairness (Procaccia 2019); envy-free allocations are fair in a very intuitive sense—each agent prefers their own allocation to any other's. In addition to being conceptually attractive, envy-freeness has the following advantages: (i) it is invariant under heterogeneous scaling (i.e, if the utility function of one agent is scaled up or down by a constant, an envy-free schedule remains envy-free). This makes it a more suitable measure of fairness than equitability and max-min fairness when the utility functions are heterogeneous; (ii) an envy-free schedule is guaranteed to exist in our setting, unlike a proportional one (we give an example of a setting where a proportional schedule does not exist in Appendix C).

Although in some cases, envy (and unfairness in general) could potentially be reduced by using monetary transfers, they are often either unethical or difficult, if not impossible, to implement; similarly to (e.g., Bertsimas et al. 2012, Gur et al. 2021), we focus on settings without auxiliary monetary transfers.

1.1. Our Contributions

We model a central planner who needs to schedule a scarce resource among several agents over some time horizon. The resource can only be allocated to a single agent at any given time. Some fixed amount of time is needed to transfer the resource between agents, during which the resource remains idle; we refer to this idle time period as the *switching cost*. We note there is a switching cost even if the resource is allocated to the same agent consecutively; e.g., a hospital room has to be disinfected when switching patients, even if both patients are treated by the same physician.

The agents are endowed with heterogeneous utilities. Each utility is a function of the length of time the resource is allocated to the agent. The agents experience temporal discounting from obtaining the resource at later times. The agents' utilities are therefore determined by the time integral of an exponentially discounted instant utility. This aligns with the literature on dynamic consumption based on Samuelson (1937) and Kahneman et al. (1997). We place a minimal requirement on the instant utility: we only require it to be non-increasing in the time the agent has held the resource. As a result, our model applies to a variety of settings where the utility functions have diminishing returns, from scheduling medical equipment such as ventilators to a parent scheduling a tablet among siblings on a flight. In particular, it captures many instant utility functions already studied in the extant dynamic consumption literature (e.g., Baucells and Sarin 2007, Baucells and Zhao 2020, Das Gupta et al. 2016). Furthermore, when there is no discounting this model recovers the set of all increasing concave utility functions as the integrals of non-negative non-increasing instant utility functions. Agents may be allocated the resource more than once, and their total utility is additive over these allocations.

As an example, consider the problem of scheduling a particle collider among five teams over the course of a month (28 days), where the switching time is two days (i.e., it takes two days between the time the previous team leaves to the time it takes the new team to start their experiment and begin obtaining their readings). As a research team spends more time with the collider, they experience diminishing marginal returns due to data saturation, technical limitations, and human fatigue. This makes their utility functions concave, as each additional unit of time yields progressively smaller benefits in data collection and experimental precision. The teams have identical utility functions: their utility as a function of time t is $u(t) = t^{0.8}$, where t is the number of days. The teams could prefer to get earlier time slots, as they would be able to publish sooner, but for simplicity, we ignore this here (this corresponds to setting the discount factor to zero in our model).

To evaluate the utility loss from requiring the schedule to be envy-free, we use the *price of fairness* (Bertsimas et al. 2011, Caragiannis et al. 2012): the ratio of the sum of the agents' utilities in an optimal utilitarian schedule and an optimal envy-free schedule (i.e., one that maximizes the total utility subject to the fairness constraint). The optimal schedule in the above example would be to allocate three of the five research teams 8 days each (and the other teams nothing), giving each of the three teams a utility of ≈ 5.28 for a total utility of ≈ 15.83. The optimal envy-free schedule would give each of the five research teams 4 days, giving each team a utility of ≈ 3.03 for a total of ≈ 15.16. The price of fairness is therefore approximately 1.044, which means that there is less than a 5% loss of efficiency as a result of requiring the schedule to be envy-free. It is easy to verify that the price of fairness can be unbounded for arbitrary problem parameters: in the above example, if the planning horizon was only 8 days long, the only envy-free schedule does not allocate any time to any team, as the total switching time needs to be at least 8 days. While the price of fairness can be unbounded in the worst-case, we take a parameterized approach, deriving bounds on the price of fairness as a function of the number of agents, the duration of the planning horizon, the switching cost, the discount factor, and the heterogeneity and concavity of the agents' utility functions. To quantify concavity, we introduce two new functionals, the *cadence* and the *discounted reciprocal function*, and show that both align with natural interpretations of concavity. The cadence aligns with the intuition that a smaller second derivative implies greater concavity. The discounted reciprocal function aligns with the idea that if two concave functions f and g intersect at the two end-points of an interval and f is above g over this interval, then f is more concave than g on this interval.

We describe efficient algorithms for computing optimal and asymptotically optimal utilitarian and envy-free schedules. For the utilitarian case, we show that when the time horizon T satisfies a certain periodic structure, a schedule in which all intervals have the same duration is optimal. Specifically, this duration is precisely the cadence of the utility of the agent with the highest "value per time unit". We then describe an algorithm for computing an envy-free schedule when the agents are homogeneous, and show that this gives an optimal schedule when the time horizon is infinite. Furthermore, we show that both of these algorithms are asymptotically optimal for all values of T.

We build upon these constructions to give three bounds on the price of fairness: an upper bound for the most general setting, and tight bounds when the time horizon is infinite and when there is no discounting. Our analysis shows that the problem parameters affect the price of fairness differently on the *macro* and *micro* levels. The macro effects dominate when the planning horizon is long, when the discount factor is large, or when the parameters change by orders of magnitude. The micro effects, in contrast, predominate when there are small changes in the parameters, the time horizon is short and the discount factor is small.

At the macro level, the price of fairness decreases in the time horizon and the degree of concavity of the agents' utility functions, and increases in the number of agents, the switching cost and the discount factor. At the micro level, however, we find that the price of fairness is not monotone in most of these parameters. This observation is particularly useful from a managerial perspective, as the central planner typically has control over some of the problem parameters. Consider again the particle collider example with the utility function $u(t) = t^{0.6}$. When there are $n = 5$ agents, the price of fairness is approximately 1.039, but when $n = 6$, the price of fairness is 1, as the optimal and envy-free allocations are the same.

Furthermore, we show that when there is no temporal discounting, problem instances exhibit *recurrent optimality* in the time horizon and the switching cost, under mild conditions. That is, if the price of fairness is not 1, it is possible to *increase* either the time horizon or the switching cost to create a problem instance with a price of fairness of 1.[2] Similarly, problem instances exhibit *asymptotic recurrent optimality* in the number of agents: it is always possible to increase the number of agents to make the price of fairness arbitrarily close to 1.

Finally, we show that when the agents are homogeneous and the number of agents is below some threshold θ, then the price of fairness is bounded by a small constant; intuitively, the threshold provides a cutoff on the number of agents that can be scheduled without sacrificing too much utility for fairness. When the agents are homogeneous and the time horizon is infinite, this threshold gives the maximal number of agents that the resource can support without sacrificing *any* utility for fairness.

2. Related Literature

Several notions of fairness are considered in the literature on resource allocation problems. Other than envy-freeness, some common notions are *proportionality*, *equitability*, and *max-min fairness.* In order to facilitate the comparison of our results to the extant literature, we briefly define these notions. An allocation is proportionally fair if each agent obtains a utility that is at least $1/n$ times

[2] Note that despite decreasing the price of fairness, increasing the switching cost can only decrease the utility of both the utilitarian and envy-free schedules.

the utility they would derive from being allocated the entire resource, an equitable allocation is one in which all agents obtain identical utilities, and a max-min fair allocation is one in which we cannot increase the utility of any agent without reducing the utility of some agent with lower utility.

2.1. The Price of Fairness

Bertsimas et al. (2011) and Caragiannis et al. (2012) independently introduced the price of fairness. Caragiannis et al. (2012) analyzed the price of proportionality, envy-freeness, and equitability in the context of dividing goods and chores; Bertsimas et al. (2011, 2012) studied fairness for resource allocation problems characterized by the utility set – the set of achievable agent utilities corresponding to all feasible resource allocations. Bertsimas et al. (2011) gave tight bounds on the price of fairness for proportionality and max-min fairness, that only depend on the number of agents n, under the assumption that the agents have equal maximum achievable utilities. Bertsimas et al. (2012) considered $\alpha-$fairness, which generalizes proportionality and max-min fairness, and characterized the trade-off with respect to a fairness parameter α. Their bounds on the price of fairness depend only on the number of agents and α, and they showed how to choose the value of α to appropriately balance efficiency and fairness. In our paper, in contrast, we assume that envy-freeness is a hard constraint, and instead of control over the fairness criterion, the planner can manipulate some of the problem parameters to reduce the value loss resulting from the fairness constraint.

The price of fairness has been studied in various operational settings. Iancu and Trichakis (2014) studied multiportfolio optimization where the fund manager would like to optimize the net utility of all accounts. Fairness considerations enter because the performance of each account depends on all the others, and the manager wants to ensure a good return on all accounts. McCoy and Lee (2014) studied maximizing the $\alpha-$fairness for a health services resource allocation problem in rural communities, subject to capacity constraints, and derived structural properties of the optimal solution. Similarly to Bertsimas et al. (2012), they studied the loss of efficiency as a function of α.

Bogomolnaia et al. (2021) considered the fair allocation of a single random object among several agents, where the allocation is constrained to be proportionally fair ex-ante. They showed that having access to the mean values of the utilities is sufficient to obtain a price of fairness that is almost as good as when the manager has access to the full distribution. Recently, Breugem and Van Wassenhove (2022) studied the price of fairness of vertical equity when it is imposed via constraints that specify a minimum percentage of the total utility for each agent, and derive an upper bound on the price of fairness.

Gur et al. (2021) studied the loss of utility as a result of provider guarantees in a centralized planning system that allocates a set of jobs to service providers. They parameterized their problem using the number of agents and the heterogeneity, and analyzed their impacts on the price of fairness. We also explicitly model the heterogeneity of the agents' utility functions; however, we allow for more general heterogeneity than Gur et al. (2021), who restricted the heterogeneity to a linear scaling coefficient. There are many other differences between our model and theirs: (i) they allocated a set of discrete, indivisible jobs, while we allocate time (which is continuously divisible); (ii) we only consider one type of constraint (envy-freeness), whereas they allowed more general constraints; and (iii) they parameterized their problem using only two parameters. Despite these differences, there are several shared insights: in the homogeneous case of Gur et al. (2021), the relative utility loss does not exceed 1/2 for any number of agents. We obtain a similar bound when the number of intervals in the optimal utilitarian schedule is at least the number of agents. In contrast to Gur et al. (2021), however, the loss incurred due to the heterogeneity in our setting is bounded by the number of agents, whereas in theirs, a large heterogeneity can lead to unbounded losses even with a small number of agents.

2.2. Envy-Free Cake-Cutting

Our model is also related to the cake-cutting framework. In this framework, a resource (the cake) is typically modeled as the interval $[0,1]$. Each agent has some non-negative integrable density function over $[0,1]$, and an agent's utility for a sub-interval is given by the integral of its density over this sub-interval. A "piece of cake" is a finite set of disjoint sub-intervals of $[0,1]$, and an agent's utility for a piece of cake is typically additive over the utilities of the sub-intervals. There is a large body of work concerning envy-free cake cutting going back to the 1950s, with a focus on the complexity of obtaining an envy-free allocation (e.g., Cohler et al. 2011, Deng et al. 2012). In addition, there have been several papers on the price of envy-freeness of cake cutting (e.g., Caragiannis et al. 2012, Aumann and Dombb 2015).

When the time horizon is finite in our model, one can represent it as the interval $[0,1]$. However, our model does not fit cleanly into the standard cake-cutting framework (even with a finite time horizon) due to the switching cost and the discounting. In the standard cake-cutting model, it is well known that there always exist proportional and envy-free allocations that are also Pareto-efficient (e.g., Weller 1985). Similarly to the standard setting, we show that there always exists an envy-free schedule. It is not difficult to verify that, as long as the time horizon is sufficiently long so as to allow each agent to be allocated some interval, any such schedule can be modified to be Pareto-efficient.

There is a related stream of literature that incorporates a temporal aspect into cake-cutting and other resource allocation problems (e.g., Walsh 2011, Lien et al. 2014, Sinclair et al. 2022, Manshadi et al. 2021). This body of work is conceptually different from our paper as the agents arrive sequentially and their demands (utility functions) are revealed upon arrival, while in our setting, the entire problem input is known upfront.

2.3. Fairness in Scheduling

Several works consider fairness in the context of scheduling, although they do not necessarily focus on the price of fairness. Moulin (2007) considered the problem of scheduling jobs among multiple agents who share a resource, with the goal of minimizing the slowdown (waiting time/job size). Qi (2017) addressed fairness in the scheduling of heterogeneous agents with uncertain service times (e.g., scheduling appointments with a physician). They introduced a type of min-max fairness and used it to schedule appointments. Breugem et al. (2021) studied fairness and attractiveness (a different notion than fairness) in crew-rostering. They devised a fairness scheme for crew-rostering, bounded its price of (max-min) fairness, and studied the trade-off between fairness and attractiveness.

Agnetis et al. (2019) and Zhang et al. (2020) consider the price of fairness in the single-machine scheduling problem with two agents. In Agnetis et al. (2019), one agent minimizes its total completion time while the other minimizes its maximum tardiness. They provide tight bounds on the price of fairness for both Kalai-Smorodinsky[3] (KS) and proportional fairness. In Zhang et al. (2020), both agents want to minimize their total completion time and one of them has exactly two jobs. They show that the price of fairness in this case is 1/2.

2.4. Temporal Utilities

Intertemporal choice models are used in many decision theory, economics, and operations research contexts to analyze how individuals make decisions involving trade-offs among costs and benefits occurring at different times. They take into account the fact that people value present rewards differently from future rewards, a concept known as time preference or discounting. The discounted utility model was first proposed by Samuelson (1937), and despite facing some criticism for its simplicity (e.g., Frederick et al. 2002), it is still the dominant model for intertemporal choice (e.g., Baucells and Sarin 2007, Bleichrodt et al. 2015). Kahneman et al. (1997) distinguished between 'decision utility' (the weight of an outcome) and 'experienced utility' (a hedonic quality), and proposed a formal model for how individuals evaluate utility over time. The primitive ingredient in this model is the 'instant utility', and the total utility is a time integral of the exponentially discounted instant utility. Kahneman et al. (1997) assumed that the instant utility is independent of the past. Their utility model is time-separable, so the same consumption leads to the same level of satisfaction at any time, independently of past consumption. Other works argue that present instant utility should depend on the history of consumption, and incorporate various forms

[3] KS fairness is equivalent to max-min fairness where all utilities are normalized.

of history-dependent instant utility. For instance, the instant utility in the models of Baucells and Sarin (2007), Baucells and Zhao (2019, 2020) is diminishing in the length/intensity of past consumption. Alternatively, the instant utility in Sundaresan (1989), Wathieu (1997), Das Gupta et al. (2016) depends on consumption relative to a history-dependent reference level, e.g., as in habit formation. Our model, where the instant utility is non-increasing in the length of time the agent has held the resource, falls under the first category.

3. Model

We consider a resource scheduling problem where a planner needs to schedule an indivisible resource among a set of n agents, over a (possibly infinite) planning horizon $0 < T \leq \infty$. Let n be the number of agents and $[n] := \{1, 2, \ldots, n\}$ be the set of agents. Let $\mathcal{U} = \{u_1, \ldots, u_n\}$ be the set of the agents' utility functions, defined below, let $\tau > 0$ be the switching cost, and let $\beta \geq 0$ be the discount factor. We denote a problem instance by $\mathcal{P} = (n, \mathcal{U}, T, \tau, \beta)$. For homogeneous utilities, we denote problem instances by $\mathcal{P}_u = (n, u, T, \tau, \beta)$, where all agents have utility function u.

3.1. Utilities

An *interval* I is a pair $I = (s, d)$ where $s \geq 0$ is the start time and $d \geq 0$ is the duration of I, respectively. Following Kahneman et al. (1997), we define, for each agent i, an *instant utility function* $\upsilon_i : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. Here, υ_i is a function of the amount of time the agent has been in possession of the resource: if agent i gets the resource at time s, the instant utility from having the resource at time $x \geq s$ is $\upsilon_i(x - s)$. To account for the temporal value loss from beginning to use the resource at a later time, we use exponential discounting (Samuelson 1937). Combining these two models (duration-dependent instant utility and exponential discounting), we define agent i's total utility from being allocated the interval $I = (s, d)$ to be:

$$u_i(I) = \int_s^{s+d} e^{-\beta x} \upsilon_i(x - s) dx, \tag{1}$$

where the instant utility depends on how long the agent has held the resource. We assume that all υ_i are non-increasing to reflect satiation from having the resource. This representation aligns with many existing dynamic consumption models, and includes several parametric forms of instant utility that have appeared in the literature: linear ($\upsilon(t) = 1 - ct$ for $c > 0$ and $t \in [0, 1/c]$) (e.g., Baucells and Zhao 2019, Das Gupta et al. 2016); power ($\upsilon(t) = t^{\alpha-1}$ for $\alpha \in (0, 1]$ and $t \geq 0$) (e.g., Baucells and Sarin 2007); and exponential ($\upsilon(t) = e^{-\alpha t}$ for $t \geq 0$) (e.g., Baucells and Sarin 2007).[4]

We overload the notation, and define $u_i(t) := u_i((0,t)) = \int_0^t e^{-\beta x} \upsilon_i(x) dx$. For $I = (s, d)$, we therefore have $u_i(I) = e^{-\beta s} u_i(d)$. We also have $u_i(0) = 0$ by Eq. (1). Since each υ_i is non-increasing, agent i's utility u_i is non-decreasing and concave as the integral of a non-negative and non-increasing function. When $\beta = 0$, we require an additional technical assumption: that $\lim_{t\to\infty} u_i(t)/t = 0$ (note that this condition is automatically satisfied for $\beta > 0$ from the definition of u_i).

An agent can be allocated multiple intervals, and their utility is additive over these intervals. Further overloading the notation, we let $u_i(\mathcal{I}) := \sum_{I \in \mathcal{I}} u_i(I)$ be agent i's utility for any set of intervals $\mathcal{I}$.

We define the heterogeneity of a set of utility functions $\mathcal{U} = \{u_1, \ldots, u_n\}$ to be the supremum of the ratios of the valuations of any pair of utility functions over their domain ($x \geq 0$):

$$\gamma(\mathcal{U}) = \sup_{u_i, u_j \in \mathcal{U}, x \geq 0} \frac{u_i(x)}{u_j(x)},$$

where $0/0 := 1$. If $\gamma(\mathcal{U}) = 1$, the utilities are homogeneous (i.e., $u_i = u_j$ for all $i, j \in [n]$).

3.2. Schedules

Let $\mathcal{I} = \{I_1, I_2, \ldots, I_m\}$ be a (possibly infinite) set of $m \geq 1$ intervals, where $I_k = (s_k, d_k)$ for $k \in [m]$. We will always assume that the intervals are sorted with respect to their start times; that is, if $k < \ell$ then $s_k < s_\ell$. We say that $\mathcal{I}$ is *feasible* for $\mathcal{P} = (n, \mathcal{U}, T, \tau, \beta)$ if we have $s_1 \geq 0$, $s_k \geq s_{k-1} + d_{k-1} + \tau$ for every $k \in [2, m]$, and if $T < \infty$ then $s_m + d_m \leq T$. The set $\mathcal{I}$ is *tight* with respect to $\mathcal{P}$ if $\mathcal{I}$ is feasible and all of the inequalities hold with equality. That is, $s_0 = 0$, $s_k = s_{k-1} + d_{k-1} + \tau$ for every $k \in [2, m]$, and if $T < \infty$ then $s_m + d_m = T$.

[4] The instant utilities in these works are obtained by first specifying a utility function (quadratic, power, or exponential) and then taking the instant utility to be its derivative.

A schedule $\pi = \{\pi_1, \ldots, \pi_n\}$ for problem $\mathcal{P}$ is an ordered set denoting the interval assignments of all agents, where π_i is the set of intervals assigned to agent i in π. Let $\mathcal{I}(\pi) = \cup_{i\in[n]}\pi_i$ be the set of all intervals assigned in π, sorted in increasing order of the interval start times. We say that π is *feasible* for $\mathcal{P} = (n, \mathcal{U}, T, \tau, \beta)$ if $\mathcal{I}(\pi)$ is feasible for $\mathcal{P}$. A schedule π is *tight* with respect to $\mathcal{P}$ if $\mathcal{I}(\pi)$ is tight with respect to $\mathcal{P}$. Note that we must have $|\mathcal{I}(\pi)| < \infty$ for $T < \infty$ because the switching cost τ is strictly positive. Let $\Pi(\mathcal{P})$ denote the set of all feasible schedules for $\mathcal{P}$.

For problem $\mathcal{P} = (n, \mathcal{U}, T, \tau, \beta)$, define the utility of agent $i \in [n]$ for schedule $\pi \in \Pi(\mathcal{P})$ to be:

$$u_i(\pi, \mathcal{P}) := \begin{cases} u_i(\pi_i), & \text{if } \beta > 0 \text{ or } T < \infty, \\ \lim_{T\to\infty} \dfrac{\sum_{k: I_k \in \pi_i, s_k + d_k \le T} u_i(I_k)}{T}, & \text{if } \beta = 0 \text{ and } T = \infty. \end{cases} \tag{2}$$

When $\beta = 0$ and $T = \infty$, we take the average utility instead of the total utility to ensure that the value is finite; this will allow us to compare the utility of two different schedules. We omit the dependence on $\mathcal{P}$ and simply write $u_i(\pi)$ when it is clear from context.

3.3. The Price of Fairness

Given problem instance $\mathcal{P}$, the unconstrained (utilitarian) scheduling problem is to maximize the sum of the agents' utilities:

$$\mathcal{P}^{UC} := \sup_{\pi \in \Pi(P)} \sum_{i\in[n]} u_i(\pi).$$

For problem $\mathcal{P}$, a schedule $\pi \in \Pi(\mathcal{P})$ is *envy-free* if $u_i(\pi_i) \ge u_i(\pi_j)$ for all $i, j \in [n]$. In other words, a schedule is envy-free if no agent prefers another agent's allocation to their own. Let $\Pi_{EF}(\mathcal{P}) \subseteq \Pi(\mathcal{P})$ denote the set of all envy-free schedules. The fairness-constrained scheduling problem is

$$\mathcal{P}^{EF} := \sup_{\pi \in \Pi_{EF}(\mathcal{P})} \sum_{i\in[n]} u_i(\pi).$$

We denote the optimal values of $\mathcal{P}^{UC}$ and $\mathcal{P}^{EF}$ by $OPT(\mathcal{P}^{UC})$ and $OPT(\mathcal{P}^{EF})$, respectively. The price of fairness for problem $\mathcal{P}$ is the ratio of the optimal total utilities in the unconstrained and fairness-constrained problems:

$$\text{PoF}(\mathcal{P}) := \frac{OPT(\mathcal{P}^{UC})}{OPT(\mathcal{P}^{EF})}.$$

It is always true that $\mathrm{PoF}(\mathcal{P}) \geq 1$, and a larger PoF indicates that a larger fraction of the total utility is lost by constraining the schedule to be envy-free.

4. Measures of Concavity

The optimal utilitarian and envy-free schedules depend on the agents' utility functions, and in particular on the degree of concavity of these functions. There is no single accepted definition of what it means for one function to be "more concave" (or more convex) than another. If for all x in some subset of the intersection of the domains of two concave functions f and g it holds that $|f''(x)| > |g''(x)|$, then it is natural to say that f is more concave than g on this set. However, if the inequality does not hold for all x in the set, it is not clear how to compare the two functions' degree of concavity. To our knowledge, the only accepted metric is that of *strong convexity*[5] (which is straightforward to adapt to concavity), which allows one to compare strongly concave functions in terms of their strong concavity parameter. However, this notion is unsuitable for our purposes, as we do not require strong concavity.

We define two novel functionals to measure the degree of concavity of a function: the *cadence* and the *discounted reciprocal function*. We show that these functionals align with natural interpretations of concavity. The cadence will be used in the construction of optimal utilitarian and envy-free schedules, and the discounted reciprocal function appears in the analysis of optimal envy-free schedules. Additionally, both functionals appear in the bounds on the price of fairness.

4.1. Cadence

We define the cadence of a function as follows.

Definition 1. For any function $f : \mathbb{R} \to \mathbb{R}$ that is non-decreasing and strictly concave on $\mathbb{R}_{\geq 0}$ and any $\tau > 0$, define

$$\textsc{Cadence}_\tau(f) := \arg\max_{t>0} \frac{f(t)}{t+\tau}. \tag{3}$$

[5] A differentiable function f is strongly convex with parameter $m \geq 0$ if $f(y) \geq f(x) + f'(x)(y-x) + (m/2)(y-x)^2$ for all $x, y \in \mathbb{R}$ (see, e.g., Vial (1982)).

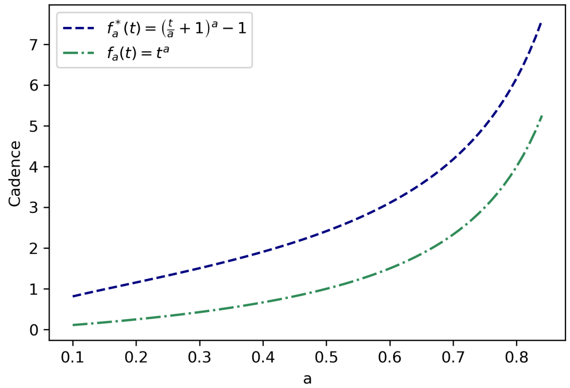

(a) Cadence_τ of $f_a(t) = t^a$ and $f_a^*(t) = \left(\frac{t}{a}+1\right)^a - 1$ as a function of a, for $\tau = 1$. The plot shows that the cadence is inversely proportional to concavity as measured by the second derivative.

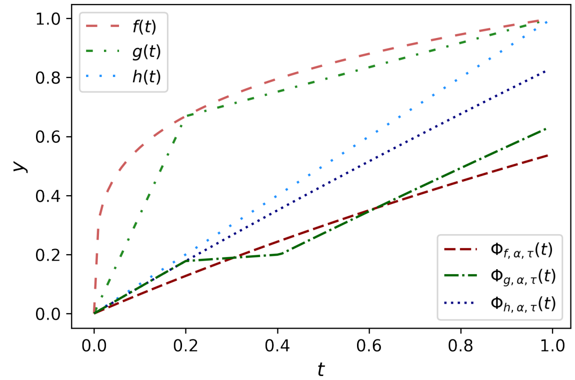

(b) Discounted reciprocal function of $f(t) = t^{0.25}$, $h(t) = t$, and g, which is a piecewise linear function that intersects f at $t = 0.2$ with parameters $\beta = 0.1$, $\tau = 1$. The functions f, g, and h intersect at $\{0, 1\}$, and f and g are more concave than h on $[0, 1]$.

Figure 1 Plots showing how the cadence and discounted reciprocal function align with natural interpretations of concavity.

In Appendix D.1, we show that the cadence is well-defined (i.e., the maximizer is finite, attained, and unique) for strictly concave functions satisfying the growth condition $\lim_{x\to\infty} u(t)/t = 0$, for any $\tau > 0$; in Appendix D.2, we show how to modify the definition of the cadence for functions that are not strictly concave.

The cadence conforms to the interpretation of the second derivative as a measure of concavity. This is exemplified in Figure 1a, which depicts the functions $f_a(t) = t^a$ and $f_a^*(t) = \left(\frac{t}{a}+1\right)^a - 1$. Both functions intuitively become more concave as a decreases, for $a \in (0, 1]$, and both $\text{Cadence}_\tau(f_a)$ and $\text{Cadence}_\tau(f_a^*)$ are increasing in a in this range. Lemma 4 in Appendix D.3 formalizes this intuition.

In the following section, we will see that the cadence corresponds to the length of the intervals in optimal utilitarian schedules. Intuitively, a function is "more concave" when its marginal utility decreases quickly, hence for such functions we would want to switch the resource more frequently.

We extend the above definition of the cadence in two ways to apply to discounted utility functions (Definition 2) and finite time horizons (Definition 3). We show in Appendix D.4 that both of these definitions indeed generalize the basic definition of cadence, by demonstrating that $\textsc{Cadence}_\tau$ is a special case thereof. First, we extend the definition to account for exponential discounting.

Definition 2. Let $f:\mathbb{R}\to\mathbb{R}$ be non-decreasing and strictly concave on $\mathbb{R}_{\geq 0}$. For any $\tau>0$ and $\beta>0$, let

$$\textsc{Cadence}_{(\beta,\tau)}(f) := \arg\max_{t>0} \frac{f(t)}{1-e^{-\beta(t+\tau)}}.$$

Definition 2 requires that $\beta>0$, but we would like the cadence to be well-defined for all $\beta\geq 0$; we therefore extend the definition of $\textsc{Cadence}_{(\beta,\tau)}$ to $\beta=0$ by setting $\textsc{Cadence}_{(0,\tau)}(f)=\textsc{Cadence}_\tau(f)$. The next definition extends the notion of cadence to finite time horizons.

Definition 3. Let $f:\mathbb{R}\to\mathbb{R}$ be non-decreasing and strictly concave on $\mathbb{R}_{\geq 0}$. For any $\tau>0$ and $T<\infty$, let

$$\textsc{Bounded-Cadence}_{(T,\tau)}(f) := \arg\max_{t:(t+\tau)\mid(T+\tau)} \frac{f(t)}{t+\tau},$$

where $(t+\tau)\mid(T+\tau)$ means that $(t+\tau)$ divides $(T+\tau)$.

We show in Appendix D.5 that all three notions of cadence can be computed efficiently for any utility function of the form of Eq. (1).

4.2. The Discounted Reciprocal Function

The discounted reciprocal function is defined as follows.

Definition 4. For any continuous and non-decreasing function $f:\mathbb{R}_{\geq 0}\to\mathbb{R}_{\geq 0}$, $\beta>0$, and $\tau>0$, let $\Phi_{f,\beta,\tau}$ be the unique value of y such that $f(y)=e^{-\beta(y+\tau)}f(t)$. We call $\Phi_{f,\beta,\tau}$ the *discounted reciprocal of f with respect to β and τ.*

In Appendix E.1, we show that the discounted reciprocal function is well-defined for all $t\geq 0$.

Intuitively, the discounted reciprocal function measures the following. Given a utility function u, we consider two subsequent intervals where the first has duration y and the second has duration t. The utility from the first interval is $u(y)$ (it is undiscounted since we get it now) and the utility from the second is $e^{-\beta(y+\tau)}u(t)$ (it is discounted since we get it starting at time $y+\tau$). Given t, $\Phi_{u,\beta,\tau}(t)$ is the required duration y of the first interval so that $u(y) = e^{-\beta(y+\tau)}u(t)$, i.e., so that both intervals have the same utility. If one agent is allocated the first interval (of length y), and another agent is allocated the second interval (of length t), then neither agent envies the other. The discounted reciprocal function is used in the computation and analysis of the price of fairness, and a variation of this function is used in the construction of the optimal envy-free schedule.

The discounted reciprocal function aligns with the following interpretation of concavity: if two concave increasing functions f and g intersect at points a and b, and for every $t \in [a,b]$, $f(t) > g(t)$, then f is more concave than g on $[a,b]$. In Figure 1b, h is linear, f and g are more concave than h on $[0,1]$ by any (reasonable) notion of concavity. Indeed $\Phi_{f,\beta,\tau}(t) \le \Phi_{h,\beta,\tau}(t)$ and $\Phi_{g,\beta,\tau}(t) \le \Phi_{h,\beta,\tau}(t)$ for all $t \in [0,1]$. We formalize this intuition in Appendix E.2. It is not clear how to compare the degree of concavity of f and g, however; f is more concave than g over $[0,0.2)$ and $(0.2,1]$, but not at $t = 0.2$. We see that neither discounted reciprocal function is greater than the other over the entire interval.

The discounted reciprocal function can be applied recursively. Define $\Phi^0_{f,\beta,\tau}(t) = t$ and for all $p \ge 1$, define the composition

$$\Phi^p_{f,\beta,\tau}(t) = \Phi_{f,\beta,\tau}\left(\Phi^{p-1}_{f,\beta,\tau}(t)\right). \tag{4}$$

We call p the *composition number* of $\Phi_{f,\beta,\tau}$.

5. Constructing Optimal and Asymptotically Optimal Schedules

In this section, we describe efficient algorithms to compute the following: (i) an optimal schedule for the unconstrained problem; and (ii) an optimal envy-free schedule when the agents are homogeneous. In the former setting, we assume that the planning horizon T is a member of a countably infinite set of real numbers, and show that this implies a construction of asymptotically optimal schedules for all values of T. For the latter setting, it is possible to obtain a similar result; however, it is not clear how to succinctly describe the values of T for which the algorithm is optimal. We therefore only prove exact optimality for $T = \infty$. In addition, we prove asymptotic optimality for general values of T. We remark upon the complexity of designing optimal envy-free schedules for heterogeneous agents in Section 7.

5.1. Optimal and Asymptotically Optimal Schedules for the Unconstrained Problem

Our first result, Theorem 1, precisely characterizes an optimal unconstrained schedule when T is a member of a countably infinite set, specifically when $T = m(\delta_{i^*} + \tau) - \tau$ for integer m. Here, δ_{i^*} is the cadence of agent i^*'s utility functions, where i^* is the agent with the highest "value per unit time" (defined formally below). The theorem states that when T is as above, there is an optimal schedule in which all intervals have the same duration: δ_{i^*}. Therefore, computing an optimal schedule can be reduced to computing the value of $\textsc{Cadence}_{(\beta,\tau)}(u_i)$ for each of the utility functions $u_i \in \mathcal{U}$. In Appendix D.5, we show that $\textsc{Cadence}_{(\beta,\tau)}(u_i)$ can be computed efficiently. Hence, the theorem yields an efficient algorithm for construction of an optimal schedule. When T is not as above, we can compute an asymptotically optimal schedule (with respect to T) as follows: choose the largest integer m such that $T_m := m(\delta + \tau) - \tau$ is not larger than T and compute the optimal schedule for T_m. More precisely, given a problem $\mathcal{P} = (n, \mathcal{U}, T, \tau, \beta)$, the schedule for $\mathcal{P}_m = (n, \mathcal{U}, T_m, \tau, \beta)$ is asymptotically optimal for $\mathcal{P} = (n, \mathcal{U}, T, \tau, \beta)$.

THEOREM 1. *Let $\mathcal{P} = (n, \mathcal{U}, T, \tau, \beta)$, $n \in \mathbb{N}_{\geq 1}$, $\tau > 0$, $\beta > 0$, and for each $i \in [n]$, denote $\delta_i = \textsc{Cadence}_{(\beta,\tau)}(u_i)$ and $\rho_i = e^{-\beta(\delta_i + \tau)}$. Set $i^* \in \arg\max_{i\in[n]} \frac{u_i(\delta_i)}{1-\rho_i}$. Let $I_k = (s_k, \delta_{i^*})$ with $s_k = (k-1)(\delta_{i^*} + \tau)$ for all k such that $s_k + \delta_{i^*} \leq T$. Set $\pi = \{\pi_1, \ldots, \pi_n\}$, where $\pi_{i^*} = \{I_k\}_{k\geq 1}$ and $\pi_j = \emptyset$ for all $j \neq i^*$. For any $m \in \mathbb{Z}$, denote $T_m = m(\delta_{i^*} + \tau) - \tau$.*

- *If $T = T_m$ for some $m \in \mathbb{Z}$, then π is an optimal schedule for $\mathcal{P}^{UC}$.*
- *Otherwise, for any $\epsilon > 0$, there exists $m_\epsilon \in \mathbb{Z}_{>0}$ such that if $T > T_{m_\epsilon}$ then*

$$\sum_{i\in[n]} u(\pi) \geq (1-\epsilon) OPT\left(\mathcal{P}^{UC}\right).$$

Proof. See Appendix F. □

5.2. Optimal and Asymptotically Optimal Envy-Free Schedules for Homogeneous Agents

We now describe an algorithm that computes an envy-free schedule π for any instance $\mathcal{P} = (n, u, T, \tau, \beta)$; the pseudocode is given as Algorithm 1. When $T = \infty$, π is an optimal EF schedule, otherwise π is asymptotically optimal. Before describing the algorithm, we need to define the following operation. Let π be a schedule for n agents. For any $t \in \mathbb{R}_{\geq 0}$, define the operation *push right* ($\triangleright$) as follows. For every interval $I_k = (s_k, d_k) \in \mathcal{I}(\pi)$, set $I'_k = (s_k + t, d_k)$. For $i \in [n]$, let $\pi'_i = \{I'_k : I_k \in \pi_i\}$, and let $\pi' = \{\pi'_1, \ldots, \pi'_n\}$. Then $\pi' = \pi \triangleright t$. Simply put, π' is the schedule where the durations of all of the intervals are identical to π, and the start times of all of the intervals are increased by t.

Algorithm 1 works as follows. Compute $\delta = \textsc{Cadence}_{(\beta,\tau)}(u)$, and let $\mathcal{I}$ be a tight set of intervals of length δ for $\mathcal{P}$. That is, $\mathcal{I} = \{I_1, I_2, \ldots\}$, where $I_k = (s_k, \delta)$ for all $k \geq 0$, $s_1 = 0$, and $s_k = s_{k-1} + \delta + \tau$ for all $k \geq 2$. Define $\rho = e^{-\beta(\delta+\tau)}$. The intervals $\mathcal{I}$ are greedily allocated among $\lambda \leq n$ agents, where $\lambda = \max_{x\in[n]} \left\{x : \rho \geq \frac{x-1}{x}\right\}$.[6] Intuitively, λ is the maximum number of agents for which such an allocation results in them all having equal total utilities. The pseudocode for this greedy algorithm is given in Appendix G. Then, the remaining agents are allocated intervals one at a time, as follows. Assume that agents $1, \ldots, i-1$ have already been allocated intervals. Allocate the interval $(0, x)$ to agent i, where $x \geq 0$ satisfies $u(x) = u(\pi_j \triangleright (x + \tau))$ for any $j \in \{1, \ldots, i-1\}$,[7] and push the schedule for the remaining agents right by $x + \tau$. This procedure guarantees that agent i's utility is identical to that of agents $1, \ldots, i-1$. Following the push right operations, the length of π is possibly greater than T. In this case, π is truncated at T to obtain π'. This may cause agents to have different utilities. Assume without loss of generality that agent n now has the lowest utility. The utility of all agents $1, \ldots, n-1$ is then decreased to $u_n(\pi')$ by shortening their intervals.

[6] We slightly abuse notation in the pseudocode, using $\textsc{Greedy}(P' = (\lambda, u, T, \tau, \beta))$ to denote assigning intervals greedily to agents $1, \ldots, \lambda$.

[7] Note that $u(\pi_j \triangleright (x + \tau))$ is identical for all $j \in \{1, \ldots, i-1\}$.

Algorithm 1: PUSH RIGHT ALGORITHM

Input: $P=(n,u,T,\tau,\beta)$

Output: An EF schedule π.

$\delta \leftarrow \text{CADENCE}_{(\beta,\tau)}(u)$

$\rho \leftarrow e^{-\beta(\delta+\tau)}$

$\lambda \leftarrow \max_{x\in[n]}\left\{x : \rho \geq \frac{x-1}{x}\right\}$

$\pi \leftarrow \text{GREEDY}(P'=(\lambda,u,T,\tau,\beta))$

for $i \in [\lambda+1, n]$ **do**

  $d_i \leftarrow y : u(y) = e^{-\beta(y+\tau)}u_{i-1}(\pi)$

  $\pi_i = (0, d_i)$

  $\pi \leftarrow \{\pi_i\} \cup \pi \triangleright (d_i+\tau)$

end

TRUNCATE π AT T

$u^{\min} \leftarrow \min_{i\in[n]}\{u_i(\pi)\}$

for $i \in [n]$ **do**

  REDUCE THE LENGTH OF THE INTERVALS IN π_i SO THAT $u_i(\pi) = u^{\min}$

end

return π

THEOREM 2. *Let $\mathcal{P}=(n,u,T,\tau,\beta)$, $n\in\mathbb{N}_{\geq 1}$, $\tau>0$, $\beta>0$, $\delta=\text{CADENCE}_{(\beta,\tau)}(u)$, and $\rho=e^{-\beta(\delta+\tau)}$. Let π be the schedule generated by Algorithm 1. Then π is envy-free, and:*

- *If $T=\infty$ then π is an optimal envy-free schedule for $\mathcal{P}$.*
- *Otherwise, for any $\epsilon>0$, there exists $T_\epsilon>0$ such that if $T>T_\epsilon$ then*

$$\sum_{i\in[n]} u(\pi) \geq (1-\epsilon) OPT\left(\mathcal{P}^{EF}\right).$$

Proof. See Appendix H. □

6. Bounds on the Price of Fairness

In this section, we present a series of theorems and propositions that establish bounds on the price of fairness under various settings and detail the dependence of these bounds on specific parameters. The theorems provide bounds on the price of fairness, applicable across different scenarios, while the propositions delve into more nuanced settings, highlighting how different parameters influence these bounds.

We prove four bounds on the price of fairness. The first, Theorem 3, is our most general bound, but it is, for the most part, not tight. This is because it is not clear how to compare the utility of the optimal envy-free schedule to that of the optimal utilitarian schedule in the most general setting. We therefore introduce several auxiliary scheduling problems as intermediaries that we can compare, to obtain a chain of inequalities starting with the optimal envy-free schedule and ending with the optimal utilitarian schedule. We lose a small factor in the bound in each inequality in exchange for generality. The other three are tight bounds for more restricted settings: Theorem 4 gives a tight bound for problem settings with an infinite time horizon and a strictly positive discount factor; Theorem 5 applies to finite horizon undiscounted settings with homogeneous utilities; and Theorem 6 is for infinite horizon undiscounted settings.

Through the propositions, we analyze the role that various parameters play in the price of fairness. These propositions are designed to complement the theorems by providing detailed insights into how the bounds are affected by various parameters unique to each scenario, and are distributed across the relevant subsections. We distinguish between *macro* and *micro* effects. Intuitively, macro effects are the large-scale effects of changing the parameters. They dominate the micro effects when, e.g., large parameter changes are made, when the time horizon is infinite, and when the discount factor β is large. Micro effects are due to small variations in the parameters; they predominate when the changes are small, the time horizon is finite, and β is small. For clarity, we focus on the case where $\beta = 0$ to highlight the micro effects, although it is straightforward to see that most of our insights still hold for sufficiently small $\beta > 0$.

6.1. A General Bound

THEOREM 3. *Let* $\mathcal{P}=(n,\mathcal{U},T,\tau,\beta)$*,* $n\in\mathbb{N}_{\geq 1}$*,* $\tau>0$*,* $\beta>0$*, and for each* $i\in[n]$*, denote* $\delta_i=\text{CADENCE}_{(\beta,\tau)}(u_i)$ *and* $\rho_i=e^{-\beta(\delta_i+\tau)}$*. Set* $i^*\in\arg\max_{i\in[n]}\frac{u_i(\delta_i)}{1-\rho_i}$*, and let* $\rho=\rho_{i^*}$*. Assume that* $T=m(\delta_{i^*}+\tau)-\tau$ *for some integer* $m>n$ *such that* $\rho^n\geq n\rho^m$*, let* $\lambda\in\mathbb{N}$ *be the largest integer such that* $\rho\geq\frac{\lambda-1}{\lambda}$*, and let* $\psi=\frac{\gamma(\mathcal{U})}{\gamma(\mathcal{U})+n-1}$*. Then if* $n>\lambda$*,*

$$\text{PoF}(\mathcal{P})\leq\frac{\lambda\psi\,(1-\rho^m)}{\rho^{n-\lambda}-\lambda\rho^m},\tag{5}$$

otherwise (if $n\leq\lambda$*),* $\text{PoF}(\mathcal{P})\leq\frac{n\psi(1-\rho^m)}{1-n\rho^m}$*.*

Proof. See Appendix I. □

Theorem 3 provides a bound on the price of fairness for problem instances with positive discounting. We restrict the time horizon T to be such that $T=m(\delta_{i^*}+\tau)-\tau$ where m is sufficiently large (similarly to Theorem 1), as we can precisely compute the optimal unconstrained schedule in this case. It is not difficult to adapt the bounds for arbitrary values of T, by extending the time horizon of the optimal unconstrained schedule to some $T'>T$, where $T'=m(\delta_{i^*}+\tau)-\tau$ for some m, as above. The following corollary to Theorem 3 covers the infinite time horizon.

COROLLARY 1. *Let* $\mathcal{P}=(n,\mathcal{U},\infty,\tau,\beta)$*,* $n\in\mathbb{N}_{\geq 1}$*,* $\tau>0$*,* $\beta>0$*, and define* ρ*,* λ*, and* ψ *as in Theorem 3. Then, if* $n>\lambda$*,*

$$\text{PoF}(\mathcal{P})\leq\frac{\lambda\psi}{\rho^{n-\lambda}},$$

otherwise $\text{PoF}(\mathcal{P})\leq n\psi$*.*

Proof. See Appendix J. □

As $\rho=e^{-\beta(\delta_{i^*}+\tau)}$, the bound of Corollary 1 suggests that the macro effects of τ and β on the price of fairness are exponential in nature. The following result, Which uses the bound of Corollary 1 and gives a matching lower bound, shows that this is indeed the case.

PROPOSITION 1. *Fix* $n\in\mathbb{N}_{\geq 2}$ *and a utility function* u*. For* $\beta>0$ *and* $\tau>0$*, let* $\mathcal{P}_{(\beta,\tau)}=(n,u,\infty,\tau,\beta)$*. Then* $\text{PoF}(\mathcal{P}_{(\beta,\tau)})$ *grows exponentially as a function of* β *and* τ*, i.e.,* $\text{PoF}(\mathcal{P}_{(\beta,\tau)})=e^{\Theta(\beta+\tau)}$*.*

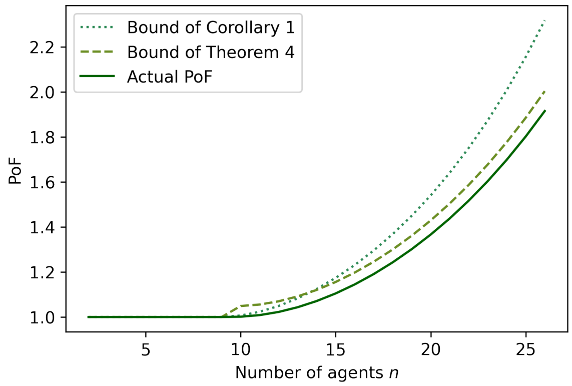

(a) PoF as a function of n.

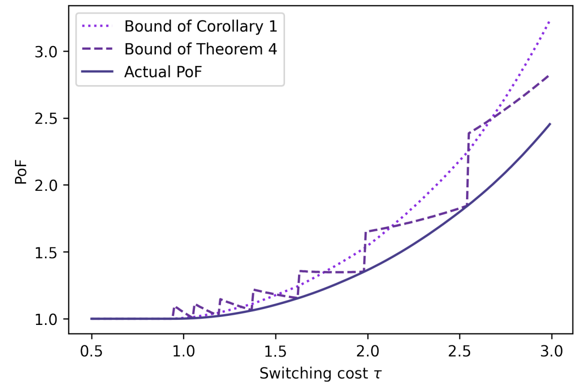

(b) PoF of $\mathcal{P}$ as a function of τ.

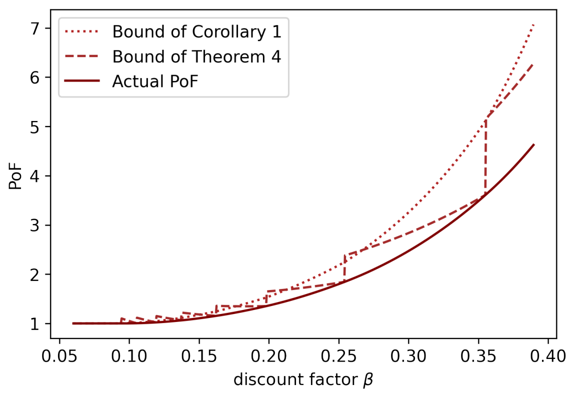

(c) PoF of $\mathcal{P}$ as a function of β.

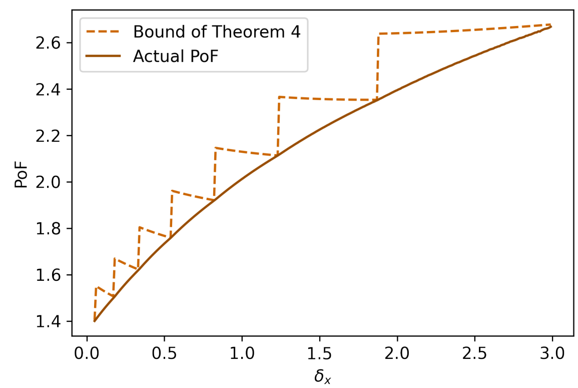

(d) PoF of $\mathcal{P} = (n = 20, u_x, \infty, \tau, \beta)$ as a function of $\delta_x := \textsc{Cadence}_{(\beta,\tau)}(u_x)$, where u_x is such that $u_x(t) = \frac{t}{x}$ for $0 \le t \le x$ and $u_x(t) = 1$ for $t > x$.

Figure 2 PoF of $\mathcal{P} = (n, u, \infty, \tau, \beta)$ as a function of n and β, where $u(t) = t^{0.1}$, $\tau = 1$, $n = 10$, and $\beta = 0.1$ unless noted otherwise.

Proof. See Appendix N.1. □

The exponential dependence of the price of fairness on β and τ described in Proposition 1 is illustrated in Figures 2b and 2c. We present the micro effects of τ in Proposition 6 below (β does not exhibit any micro effects). The macro effects of n are discussed in the next subsection, as the proof will rely on Theorem 4.

6.2. A Tight Bound for the Infinite Horizon Problem

Theorem 4 gives tight bounds on the price of fairness for problems with an infinite time horizon.

Theorem 4. *Let* $\mathcal{P} = (n,\mathcal{U},\infty,\tau,\beta)$*,* $n \in \mathbb{N}_{\geq 1}$*,* $\tau > 0$*,* $\beta > 0$*, and for each* $i \in [n]$*, denote* $\delta_i = \textsc{Cadence}_{(\beta,\tau)}(u_i)$ *and* $\rho_i = e^{-\beta(\delta_i+\tau)}$*. Set* $i^* \in \arg\max_{i\in[n]} \frac{u_i(\delta_i)}{1-\rho_i}$*, and let* $\rho = \rho_{i^*}$*. Let* $\lambda \in \mathbb{N}$ *be the largest integer such that* $\rho \geq \frac{\lambda-1}{\lambda}$*, and let* $\psi = \frac{\gamma(\mathcal{U})}{\gamma(\mathcal{U})+n-1}$*. If* $\lambda \geq n$ *then* $\mathrm{PoF}(\mathcal{P}) \leq n\psi$*, otherwise*

$$\mathrm{PoF}(\mathcal{P}) \leq \frac{\psi u_{i^*}(\delta_{i^*})}{(1-\rho)\, u_{i^*}\left(\Phi^{n-\lambda}_{u_{i^*},\beta,\tau}(\delta_{i^*})\right)}. \tag{6}$$

Furthermore, there exists a family of instances for which this bound is tight.

Proof. See Appendix K. □

The term $\rho^{n-\lambda}$ in the denominator of Inequality (5) in Theorem 3 suggests that the dependence on n is exponential, as λ is independent of n. The same dependence appears in Corollary 1. However, Theorem 3 and Corollary 1 are not tight. To show that the dependence is exponential when the time horizon is infinite, we have the following proposition, which relies on a deeper analysis of the bound of Theorem 4. We note that when the time horizon is finite, the dependence on n is no longer exponential (see Proposition 5 below).

Proposition 2. *For any* $\tau > 0$*,* $\beta > 0$*, and* $\lambda \in \mathbb{N}_{\geq 1}$*, there exists a family of instances* $\{\mathcal{P}_n\}_{n=\lambda}^{\infty}$ *where* $\mathcal{P}_n = (n,u,\infty,\tau,\beta)$ *such that for* $n = \lambda$*,* $\mathrm{PoF}(\mathcal{P}_n) = 1$*, and for* $n > \lambda$*, the price of fairness grows exponentially in* n*, i.e.,* $\mathrm{PoF}(\mathcal{P}_n) = \Omega(2^n)$*.*

Proof. See Appendix N.2. □

We now analyze the macro effects of the heterogeneity of the utility functions. In the following proposition, we show that the heterogeneity cannot affect the price of fairness by a factor greater than n. To formalize this statement, let $\mathcal{P}_u = (n,u,\infty,\tau,\beta)$ be a problem instance where u is arbitrary, and consider any $\mathcal{P} = (n,\mathcal{U},\infty,\tau,\beta)$, where $u_1 = u$ and $u_2,\ldots,u_n$ are upper bounded by u_1 and lower bounded by u_1/γ for some $\gamma \geq 1$. The heterogeneity of $\mathcal{U}$ is clearly at most γ, and we show that for any γ, the price of fairness of $\mathcal{P}$ is at most n times the price of fairness of $\mathcal{P}_u$.

Proposition 3. *Let* $\mathcal{P}_u = (n,u,\infty,\tau,\beta)$*,* $n \in \mathbb{N}_{\geq 1}$*,* $\tau > 0$*,* $\beta \geq 0$*, for arbitrary* u*. Let* $\gamma \in \mathbb{R}_{\geq 1}$*, and let* $\mathcal{P} = (n,\mathcal{U},\infty,\tau,\beta)$*, where* $\mathcal{U} = \{u_1,\ldots,u_n\}$*,* $u_1 = u$*, and* u_i *satisfies* $\frac{u(t)}{\gamma} \leq u_i(t) \leq u(t)$ *for all* $i \in [2,n]$ *and* $t \geq 0$*. Then*

$$\mathrm{PoF}(\mathcal{P}) \leq \min\{\gamma,n\} \cdot \mathrm{PoF}(\mathcal{P}_u).$$

Proof. See Appendix N.3. □

It is more difficult to quantify the effect of the concavity of the utility functions on the price of fairness, in part since there is not a unique measure of concavity. The dependence on the concavity appears in Theorem 3 and Corollary 1 through the cadence only; however these bounds are not tight. In the tight bound in Theorem 4, the dependence on the concavity appears through both the cadence and the discounted reciprocal function. To simplify the effect of the concavity, we consider a family of problem instances that have identical cadences, so that the difference in concavity is measured only through the discounted reciprocal function.

INSTANCE 1. *Fix $\beta = 0.1$, let $\tau > 0$ be such that $e^{-\beta(\tau+1)} = \frac{3}{4}$, and let $n \in \mathbb{N}_{\geq 4}$. For any $x \in [0.5, 1]$, set $u_x(t) = t^x$ for $t \in [0,1]$ and $u_x(t) = 1$ for $t > 1$. For $x \in [0.5, 1]$, denote $\mathcal{P}_x = (n, u_x, \infty, \tau, \beta)$.*

The following result shows that a higher concavity leads to a lower price of fairness.

PROPOSITION 4. *Let β, τ, and n be as in Instance 1. Then for every $x, x' \in [0.5, 1]$ such that $x < x'$, it holds that* $\mathrm{PoF}(\mathcal{P}_x) < \mathrm{PoF}(\mathcal{P}_{x'})$.

Proof. See Appendix N.4. □

The same qualitative dependence on the concavity (i.e., that the price of fairness is inversely correlated with the concavity) appears in Corollary 1. Quantitatively, the bound of Corollary 1 increases exponentially with the cadence. Figure 2d shows that this bound is far from tight, and that not only is the true dependence on the cadence not exponential, it is not even convex.

6.3. Tight Bounds for the Undiscounted Problem

The following result gives tight upper bounds on the price of fairness when $\beta = 0$.

THEOREM 5. *Let $\mathcal{P} = (n, u, T, \tau, 0)$, $n \in \mathbb{N}_{\geq 1}$, $0 < T < \infty$, and $\tau > 0$. Let $\delta =$ BOUNDED-CADENCE$_{(T,\tau)}(u)$, and $m = \frac{T+\tau}{\delta+\tau}$. Then*

$$\mathrm{PoF}(\mathcal{P}) \leq \begin{cases} \frac{m}{m-(m \mod n)}, & \text{if } n \leq m, \\ \frac{T-(m-1)\tau}{T-(n-1)\tau}, & \text{if } n > m \text{ and } T > (n-1)\tau. \end{cases}$$

Otherwise, if $T \le (n-1)\tau$*, then* $\mathrm{PoF}(\mathcal{P}) = \infty$*.*

Furthermore, for any $\tau > 0$ *and* $n \in \mathbb{N}_{\ge 1}$ *there exists a function* u *such that the above inequalities hold with equality for* $\mathcal{P}_T = (n, u, T, \tau, 0)$ *for countably many values of* T*.*

Proof. See Appendix L. □

The parameter m in the theorem statement corresponds to the number of intervals in the optimal unconstrained schedule. Note that by the definition of BOUNDED-CADENCE, m in the theorem statement is an integer. When $T < \infty$ and $\beta = 0$, the macro effects of n are as follows. There is some threshold θ such that if $n \le \theta$, the price of fairness is bounded by a constant, and when $n > \theta$ it is increasing and convex[8] in n.

PROPOSITION 5. *For any* $\tau > 0$*,* $T < \infty$*, and utility function* u*, let* $\mathcal{P}_n = (n, u, T, \tau, 0)$ *for* $n \in \mathbb{N}_{\ge 2}$*. There exists some* $m < \frac{T}{\tau} + 1$ *such that* $\mathrm{PoF}(\mathcal{P}_n) \le 2$ *if* $n < m$*,* $\mathrm{PoF}(\mathcal{P}_n) = \infty$ *if* $n \ge \frac{T}{\tau} + 1$*, and* $\mathrm{PoF}(\mathcal{P}_n)$ *is increasing and convex (in the discrete sense) in* n *if* $m \le n < \frac{T}{\tau} + 1$*.*

Proof. See Appendix N.5. □

The macro effects of τ when $\beta = 0$ can be seen in Figure 3c. From Theorem 5, we can see that the price of fairness is actually unbounded. As τ grows, the number of intervals in the optimal utilitarian schedule decreases, and as $\tau \to \frac{T}{n-1}$, the denominator of $\frac{T-(m-1)\tau}{T-(n-1)\tau}$ approaches zero.

When $\beta = 0$ and $T < \infty$, we can see the micro-effects come into play: a little flexibility in the number of participating agents, the length of the time horizon, or the switching cost allows us to reduce the utility loss resulting from the fairness requirement to zero or almost zero. To formalize this concept, we use the following definition.

DEFINITION 5. Let u be a utility function, and for any $n \in \mathbb{N}_{>0}$ and $0 \le \tau \le T \le \infty$, denote $\mathcal{P}_{n,\tau,T} = (n, u, T, \tau, 0)$. The set of instances $\{\mathcal{P}_{n,\tau,T} | n \in \mathbb{N}_{>0}, 0 \le \tau \le T \le \infty\}$ is

1. *Recurrently Optimal in* T if for any $n > 0$, $\tau > 0$, and integer $c > 0$, there exists a $T_c > 0$ and a sequence $0 < T_1 < T_2 < \cdots < T_c$ such that $\mathrm{PoF}(\mathcal{P}_{n,\tau,T_i}) = 1$ for all $i \in [c]$.

[8] Convex in the discrete sense.

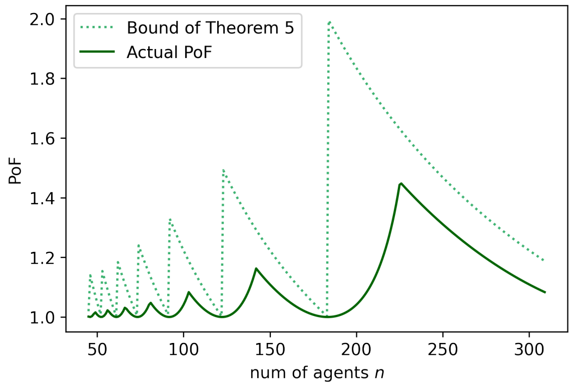

(a) PoF as a function of the number of agents n, where $u(t) = 1 - e^{-3t}, T = 1003, \tau = 2$.

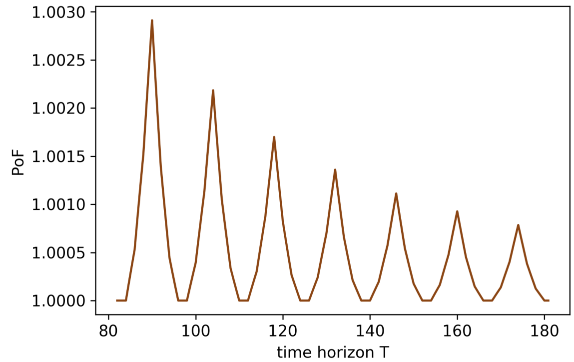

(b) PoF as a function of the time horizon T, where $u(t) = \sqrt{t}, \tau = 1, n = 7$.

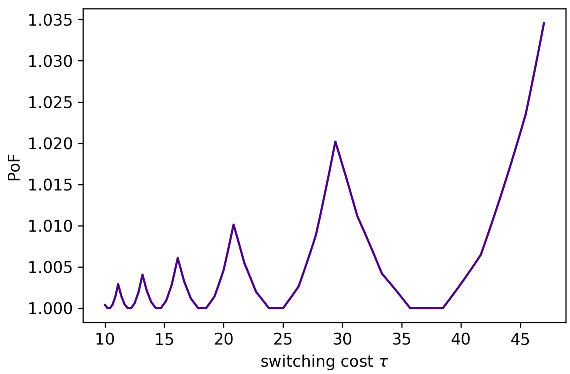

(c) PoF as a function of the switching cost τ, where $u(t) = \sqrt{t}, T = 1000, n = 7$.

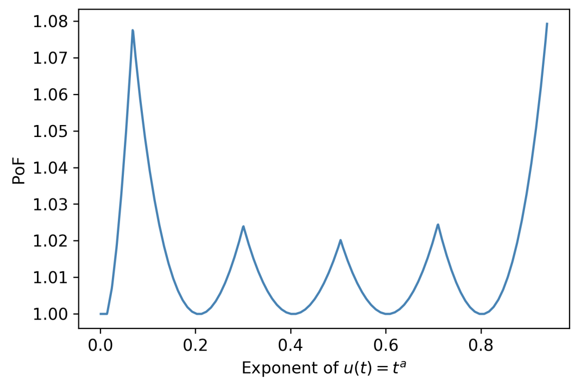

(d) PoF as a function of the concavity of the utility function $u(t) = t^a$, parameterized by a, with $T = 100, \tau = 1, n = 20$.

Figure 3 The price of fairness as a function of the time horizon, switching cost, number of agents and concavity of the utility function. For all sub-figures, $\beta = 0$.

2. *Recurrently Optimal in* τ if for any $n > 0$ and integer $c > 0$, there is a $T > 0$ and a sequence $0 < \tau_1 < \tau_2 < \cdots < \tau_c$ such that $\text{PoF}(\mathcal{P}_{n,\tau_i,T}) = 1$ for all $i \in [c]$.
3. *Asymptotically Recurrently Optimal in* n if for any $\tau > 0$, $\epsilon > 0$, and integer $c > 0$, there is a $T > 0$ and a sequence $0 < n_1 < n_2 < \cdots < n_c$ such that $\text{PoF}(\mathcal{P}_{n_i,\tau,T}) < 1 + \epsilon$ for all $i \in [c]$.

Proposition 6. *For any utility function* u*, the set of instances* $\{\mathcal{P}_{n,\tau,T} = (n, u, T, \tau, 0), n \in \mathbb{N}_{>0}, 0 \leq \tau \leq T \leq \infty\}$ *is recurrently optimal in* τ *and* T *and asymptotically recurrently optimal in* n.

Proof. See Appendix N.6. □

Figures 3a, 3b, and 3c illustrate the results of Proposition 6 with respect to n, T, and τ, respectively. We chose $T = 1003$ (and not, e.g., $T = 1000$) in Figure 3a to illustrate that the set of instances is only asymptotically recurrently optimal. For this value of T, the number of intervals in the optimal schedule is 367, a prime number. Therefore, there is no value of n (other than $n \in \{1, 367\}$) for which the price of fairness is one. Nevertheless, we can reduce the value loss significantly by appropriately selecting the number of agents. Note that while increasing τ may reduce the price of fairness, the utilities of both the optimal utilitarian and the optimal fair schedule are non-increasing in τ.

For completeness, we include a bound for the infinite horizon undiscounted case.

THEOREM 6. *Let* $\mathcal{P} = (n, \mathcal{U}, \infty, \tau, 0)$*,* $n \in \mathbb{N}_{\geq 1}$*,* $\tau > 0$*, and* $\psi = \frac{\gamma(\mathcal{U})}{\gamma(\mathcal{U})+n-1}$*. Then* $\mathrm{PoF}(\mathcal{P}) \leq n\psi$*. Furthermore, for any* $n \in \mathbb{N}_{\geq 1}$ *and* $\tau > 0$*, there exists a set of* n *utility functions* $\mathcal{U}$ *for which this bound is tight.*

7. Conclusion, Limitations, and Future Research

In this paper, we analyzed the contribution of the different problem parameters to the loss of utility resulting from constraining the schedules to be envy-free. We found that the parameters affect the price of fairness differently on the macro and micro levels, and derived insights into how a schedule planner could reduce the price of fairness by tuning the parameters over which they have control.

Our results lay the groundwork for further exploration in this area. First, our model has several simplifying assumptions, and it would be interesting to remove some (or all) of them in future research. For instance, we assume the switching cost is fixed, the discount factor is uniform, and the agents' utilities only depend on the start times and durations of their assigned intervals. None of these might be the case in practice. Second, while envy-freeness is an attractive notion of fairness, one could study the price of fairness in resource scheduling for other fairness notions. Third, it may be interesting to study dynamic variations, where part of the problem input is uncertain and is revealed in an online fashion. There are several ways in which uncertainty could be incorporated into the model: the agents' utilities could change as new information becomes available; the switching cost for a medical device could be stochastic; the number of agents and time horizon could change dynamically.

Finally, we remark on the complexity of computing optimal envy-free schedules for heterogeneous agents. Algorithm 1, which computes an asymptotically optimal envy-free schedule for problem instances with homogeneous agents, can be implemented efficiently. In the proof of Theorem 3 (in the Appendix), we demonstrate the existence of an envy-free schedule for heterogeneous agents using the Brouwer fixed point theorem. This theorem is not constructive, and computing a fixed point is known to be PPAD-hard (Hirsch et al. 1989, Chen and Deng 2009). We conjecture that computing an envy-free schedule for heterogeneous agents is NP-hard. In future work, it would be interesting to explore tractable approximation schemes for envy-free schedules for heterogeneous agents.

References

Agnetis A, Chen B, Nicosia G, Pacifici A (2019) Price of fairness in two-agent single-machine scheduling problems. *European Journal of Operational Research* 276(1):79–87, ISSN 0377-2217.

Aumann Y, Dombb Y (2015) The efficiency of fair division with connected pieces. *ACM Trans. Econ. Comput.* 3(4), ISSN 2167-8375.

Baucells M, Sarin RK (2007) Satiation in discounted utility. *Operations Research* 55(1):170–181.

Baucells M, Zhao L (2019) It is time to get some rest. *Management Science* 65(4):1717–1734.

Baucells M, Zhao L (2020) Everything in moderation: Foundations and applications of the satiation model. *Management Science* 66(12):5701–5719.

Bertsimas D, Farias VF, Trichakis N (2011) The price of fairness. *Operations Research* 59(1):17–31.

Bertsimas D, Farias VF, Trichakis N (2012) On the efficiency-fairness trade-off. *Management Science* 58(12):2234–2250.

Bleichrodt H, Keskin U, Rohde KIM, Spinu V, Wakker P (2015) Discounted utility and present value—a close relation. *Operations Research* 63(6):1420–1430.

Bogomolnaia A, Moulin H, Sandomirskiy F (2021) On the fair division of a random object. *Management Science* .

Breugem T, Dollevoet T, Huisman D (2021) Is equality always desirable? analyzing the trade-off between fairness and attractiveness in crew rostering. *Management Science* .

Breugem T, Van Wassenhove LN (2022) The price of imposing vertical equity through asymmetric outcome constraints. *Management Science* 0(0):null.

Caragiannis I, Kaklamanis C, Kanellopoulos P, Kyropoulou M (2012) The efficiency of fair division. *Theory of Computing Systems* 50(4):589–610, ISSN 1433-0490.

Chen X, Deng X (2009) On the complexity of 2d discrete fixed point problem. *Theoretical Computer Science* 410(44):4448–4456, ISSN 0304-3975, automata, Languages and Programming (ICALP 2006).

Cohler YJ, Lai JK, Parkes DC, Procaccia AD (2011) Optimal envy-free cake cutting. *Twenty-Fifth AAAI Conference on Artificial Intelligence*.

Das Gupta A, Karmarkar US, Roels G (2016) The design of experiential services with acclimation and memory decay: Optimal sequence and duration. *Management Science* 62(5):1278–1296.

Deng X, Qi Q, Saberi A (2012) Algorithmic solutions for envy-free cake cutting. *Operations Research* 60(6):1461–1476.

Frederick S, Loewenstein G, O'Donoghue T (2002) Time discounting and time preference: A critical review. *Journal of Economic Literature* 40(2):351–401.

Gur Y, Iancu D, Warnes X (2021) Value loss in allocation systems with provider guarantees. *Management Science* 67(6):3757–3784.

Hirsch MD, Papadimitriou CH, Vavasis SA (1989) Exponential lower bounds for finding brouwer fix points. *Journal of Complexity* 5(4):379–416, ISSN 0885-064X.

Iancu DA, Trichakis N (2014) Fairness and efficiency in multiportfolio optimization. *Operations Research* 62(6):1285–1301.

Johnston M (1990) Spike: AI scheduling for NASA's Hubble Space Telescope. *Sixth Conference on Artificial Intelligence for Applications*, 184,185,186,187,188,189,190 (Los Alamitos, CA, USA: IEEE Computer Society).

Kahneman D, Wakker PP, Sarin R (1997) Back to bentham? explorations of experienced utility. *The quarterly journal of economics* 112(2):375–406.

Lien RW, Iravani SMR, Smilowitz KR (2014) Sequential resource allocation for nonprofit operations. *Operations Research* 62(2):301–317.

Manshadi VH, Niazadeh R, Rodilitz S (2021) Fair dynamic rationing. Biró P, Chawla S, Echenique F, eds., *EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021*, 694–695 (ACM).

McCoy JH, Lee HL (2014) Using fairness models to improve equity in health delivery fleet management. *Production and Operations Management* 23(6):965–977.

Moulin H (2007) Minimizing the worst slowdown: Offline, online. *Operations Research* 55(5):876–889.

Procaccia AD (2019) Axioms should explain solutions. *The Future of Economic Design*, 195–199 (Springer).

Qi J (2017) Mitigating delays and unfairness in appointment systems. *Management Science* 63(2):566–583.

Samuelson PA (1937) A note on measurement of utility. *The Review of Economic Studies* 4(2):155–161, ISSN 00346527, 1467937X.

Shetrone M, Cornell ME, Fowler JR, Gaffney N, Laws B, Mader J, Mason C, Odewahn S, Roman B, Rostopchin S, Schneider DP, Umbarger J, Westfall A (2007) Ten year review of queue scheduling of the Hobby-Eberly Telescope. *Publications of the Astronomical Society of the Pacific* 119(855):556–566.

Sinclair SR, Banerjee S, Yu CL (2022) Sequential fair allocation: Achieving the optimal envy-efficiency tradeoff curve. Manjunath D, Nair J, Carlsson N, Cohen E, Robert P, eds., *SIGMETRICS/PERFORMANCE '22: ACM SIGMETRICS/IFIP PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems, Mumbai, India, June 6 - 10, 2022*, 95–96 (ACM).

Sundaresan SM (1989) Intertemporally dependent preferences and the volatility of consumption and wealth. *The Review of Financial Studies* 2(1):73–89, ISSN 08939454, 14657368.

Vial JP (1982) Strong convexity of sets and functions. *Journal of Mathematical Economics* 9(1):187–205, ISSN 0304-4068.

Walsh T (2011) Online cake cutting. *Algorithmic Decision Theory* 292–305.

Wathieu L (1997) Habits and the anomalies in intertemporal choice. *Management Science* 43(11):1552–1563, ISSN 00251909, 15265501.

Weller D (1985) Fair division of a measurable space. *Journal of Mathematical Economics* 14(1):5–17.

Zhang Y, Zhang Z, Liu Z (2020) The price of fairness for a two-agent scheduling game minimizing total completion time. *Journal of Combinatorial Optimization* 44(3), ISSN 1573-2886.

Appendix A: Summary of Notation

Symbol	Description	Definition	Remarks
T	Time horizon	-	$T > 0$
τ	Switching time	-	$\tau > 0$
β	Discount factor	-	$\beta \geq 0$
$\gamma(\mathcal{U})$	Heterogeneity	$\sup_{u_i, u_j \in U, x \geq 0} \frac{u_i(x)}{u_j(x)}$	$\gamma(\mathcal{U}) \geq 1$
ψ	-	$\frac{\gamma(\mathcal{U})}{\gamma(\mathcal{U})+n-1}$	-
$\text{CADENCE}_\tau(f)$	Cadence	$\arg\max_{t>0} \frac{f(t)}{t+\tau}$	-
$\text{CADENCE}_{(\beta,\tau)}(u)$	Cadence	$\arg\max_{t>0} \frac{u(t)}{1-e^{-\beta(t+\tau)}}$	Notated by δ
$\text{BOUNDED-CADENCE}_{(T,\tau)}(u)$	Bounded Cadence	$\arg\max_{t:(t+\tau)\mid(T+\tau)} \frac{f(t)}{t+\tau}$	Notated by δ
ρ_i	-	$e^{-\beta(\delta_i+\tau)}$	-
i^*	'Optimal agent'	$\in \arg\max_{i\in[n]} \frac{u_i(\delta_i)}{1-\rho_i}$	-
λ	-	$\max x : \rho_{i^*} \geq \frac{x-1}{x}$	Integer
$\Phi_{f,\beta,\tau}(t)$	Discounted Reciprocal	$y \geq 0 : f(y) = e^{-\beta(y+\tau)} f(t)$	-
$OPT(\mathcal{P}^{UC})$	Value of optimal utilitarian schedule for $\mathcal{P}$		
$OPT(\mathcal{P}^{EF})$	Value of optimal envy-free schedule for $\mathcal{P}$		

Table 1 Table of Notation

Appendix B: Results on Partitions of Geometric Series

We prove several results on partitions of geometric series $\{a, ar, ar^2, \ldots, ar^z\}$ for $z \in \mathbb{N}_{\geq 0}$. Lemma 1 assumes that r is sufficiently large, and Lemma 2 extends the result of Lemma 1 to more general values of r.

LEMMA 1. *Fix $n \in \mathbb{N}_{>0}, a > 0$. If $r \in [\frac{n-1}{n}, 1)$, then for every $z \in \mathbb{N}_{\geq 0}$ there exists a partition $A_1, \ldots, A_n$ of $\{a, ar, ar^2, \ldots, ar^z\}$ such that:*

(i) For all $j \in [n]$, $\sum_{x \in A_j} x \in \left[\frac{a(1-nr^{z+1})}{n(1-r)}, \frac{a}{n(1-r)}\right]$.

(ii) There exists some $j \in [n]$ for which $\sum_{x \in A_j} x \leq \frac{a}{n(1-r)} - ar^{z+1}$.

Furthermore, this partition can be found using a greedy assignment.

Proof. Recall that $\sum_{i=0}^{z} ar^i = \frac{a(1-r^{z+1})}{1-r}$. Note that $0 \leq n(1-r) < 1$, hence $\frac{1}{n(1-r)} > 1$. Fix n and r. The proof is by induction on z.

Base case. To show (i), note that for $z = 0$, $\sum_{x \in A_j} x \in \{0, a\}$ for all j, hence the upper bound of (i) holds. To see the lower bound of (i), note that

$$\frac{a\left(1-nr^{z+1}\right)}{n(1-r)} = \frac{a\left(1-nr\right)}{n(1-r)}. \tag{7}$$

If $n = 1$, (7) equals a, $\sum_{x \in A_1} x = a$, and the lower bound holds. If $n \geq 2$, $1 - nr \leq 0$ (as $r \geq \frac{1}{2}$), hence (7) is non-positive; as $\sum_{x \in A_j} x \geq 0$, the bound holds.

If $n = 1$ then (ii) holds as $\frac{1}{1-r} - r \geq 1$ (with equality if $r = 0$). If $n > 1$ then (ii) holds as there is some $j \in [n]$ for which $\sum_{x \in A_j} x = 0 < a - ar^{z+1} < \frac{a}{n(1-r)} - ar^{z+1}$.

Inductive step. Assume that the lemma holds for z. Denote the partition of $\{a, ar, \ldots, ar^z\}$ by $A_1^z, \ldots, A_n^z$ and take $j' = \arg\min_j \sum_{x \in A_j^z} x$. Set the partitions of $\{a, ar, \ldots, ar^z, ar^{z+1}\}$ to be $A_j^{z+1} = A_j^z$ for all $j \neq j'$, and $A_{j'}^{z+1} = A_{j'}^z \cup \{ar^{z+1}\}$. We prove the three bounds: the upper and lower bounds for (i) and the (upper) bound of (ii):

1. Upper bound of (i): For all $j \neq j'$, $\sum_{x \in A_j^{m+1}} x \leq \frac{a}{n(1-r)}$ from the induction hypothesis (part (i)). It must hold that $\sum_{x \in A_{j'}^z} x \leq \frac{a}{n(1-r)} - ar^{z+1}$ from the induction hypothesis (part (ii)), hence $\sum_{x \in A_{j'}^{z+1}} x \leq \frac{a}{n(1-r)}$.
2. Lower bound of (i): If there exists j'' such that $\sum_{x \in A_{j''}} x < \frac{a(1-nr^{z+2})}{n(1-r)}$, then

$$\begin{aligned}\sum_{j \neq j''} \sum_{x \in A_j} x &= \sum_j \sum_{x \in A_j} x - \sum_{x \in A_{j''}} x \\ &> \frac{a\left(1-r^{z+2}\right)}{1-r} - \frac{a\left(1-nr^{z+2}\right)}{n(1-r)} \\ &= \frac{a(n-1)}{n(1-r)},\end{aligned} \tag{8}$$

where Ineq. (8) is due to the assumption that $\sum_{x \in A_{j''}} x < \frac{a(1-nr^{z+2})}{n(1-r)}$. However, then there is at least one j for which $\sum_{x \in A_j} x$ violates the upper bound, as the upper bound implies that $\sum_{j \neq j''} \sum_{x \in A_j} x \leq (n-1)\frac{a}{n(1-r)}$.

3. Part (ii). If for all $j \in [n]$, it holds that $\sum_{x\in A_j} x > \frac{a}{n(1-r)} - ar^{z+2}$ then

$$\sum_{j\in[n]}\sum_{x\in A_j} x > \frac{an}{n(1-r)} - anr^{z+2} \geq \frac{a\left(n - nr^{z+2}\right)}{n(1-r)},$$

where the second inequality is because $\frac{1}{n(1-r)} \geq 1$. This is in contradiction to the fact that $\sum_{x\in A_j} x = \frac{a\left(1-r^{z+2}\right)}{1-r}$. □

Taking the limits in part (i) of Lemma 1, the following corollary is immediate.

COROLLARY 2. *Fix $n \in \mathbb{N}_{>0}$ and $a > 0$. If $r \in [1-\frac{1}{n}, 1)$, then there exists a partition $A_1, \ldots, A_n$ of $\{a, ar, ar^2, \ldots\}$ such that for all $j \in [n]$, $\sum_{x\in A_j} x = \frac{a}{n(1-r)}$. Furthermore, this partition can be found using a greedy assignment.*

LEMMA 2. *Let $n \in \mathbb{N}_{\geq 1}$, $a > 0$, $r \in [0, 1)$, and $\mathcal{A} = \{a, ar, ar^2, \ldots, ar^z\}$ for some $z \in \mathbb{N}_{\geq 0}$. Let $\lambda \in \mathbb{N}$ be the largest integer such that $r \geq \frac{\lambda-1}{\lambda}$. Then:*

(i) If $n \leq \lambda$, there exists a partition $A_1, \ldots, A_n$ of $\mathcal{A}$ such that $\displaystyle\sum_{x\in A_j} x \geq \frac{a\left(1 - nr^{z+1}\right)}{n(1-r)}$ for every $j \in [n]$.

(ii) If $n > \lambda$, there exists a partition $A_1, \ldots, A_n$ of $\mathcal{A}$ such that $\displaystyle\sum_{x\in A_j} x \geq \frac{a\left(r^{n-\lambda} - \lambda r^{z+1}\right)}{\lambda(1-r)}$ for every $j \in [n]$.

Proof. If $n \leq \lambda$ then $r \in \left[\frac{n-1}{n}, 1\right)$, and the lemma follows immediately from Lemma 1. Otherwise ($r < \frac{n-1}{n}$), set $A_i = \{ar^{i-1}\}$ for agents $i \in [1, n-\lambda]$. Then, for all $i \in [1, n-\lambda]$ it is true that

$$\begin{aligned} ar^{i-1} &\geq ar^{n-\lambda-1} \\ &\geq \frac{ar \cdot r^{n-\lambda-1}}{\lambda(1-r)} \qquad (9) \\ &= \frac{ar^{n-\lambda}}{\lambda(1-r)} \\ &\geq \frac{a\left(r^{n-\lambda} - \lambda r^{z+1}\right)}{\lambda(1-r)}, \end{aligned}$$

where Ineq. (9) follows because $r \leq \frac{\lambda}{\lambda+1}$, and therefore $\frac{r}{\lambda(1-r)} \leq 1$.

It remains to allocate the set $\mathcal{A}^{\ddagger} = \{ar^{n-\lambda}, ar^{n-\lambda+1}, \ldots, ar^z\}$ among the λ remaining agents. Instead, we allocate the infinite set $\{ar^{n-\lambda}, ar^{n-\lambda+1}, \ldots\}$, noting that we will have allocated the elements $ar^{z+1}\, ar^{z+2}, \ldots$, which are not in $\mathcal{A}^{\ddagger}$. Substituting $n^{\ddagger} = \lambda$ and $a^{\ddagger} = ar^{n-\lambda}$, we would like to allocate $a^{\ddagger}, a^{\ddagger}r, \ldots$ among $n^{\ddagger}$ remaining agents, where $r \in [1 - \frac{1}{n^{\ddagger}}, 1)$. The conditions of Corollary 2 hold for $n^{\ddagger}, r^{\ddagger}$, and there exists a partition $A^{\ddagger}_{n-\lambda+1}, \ldots, A^{\ddagger}_n$ such that for all $j \in [n-\lambda+1, n]$,

$$\sum_{x\in A^{\ddagger}_j} x = \frac{a^{\ddagger}}{n^{\ddagger}(1-r)} = \frac{ar^{n-\lambda}}{\lambda(1-r)}.$$

Recall that our allocation includes the set $\{ar^{z+1}, ar^{z+2}, \ldots\}$. It holds that

$$\sum_{k=z+1}^{\infty} ar^k = \frac{ar^{z+1}}{1-r}.$$

Therefore

$$\sum_{x \in A_j} x = \frac{a\left(r^{n-\lambda} - \lambda r^{z+1}\right)}{\lambda(1-r)},$$

completing the proof. □

Appendix C: Additional Material for Section 1

C.1. A Proportionally Fair Schedule Might Not Exist

Example 1. Consider a problem with time horizon $T = 16$, switching cost $\tau = 14$, two agents with utility functions $u(t) = \sqrt{t}$, and no discounting. If a single agent is allocated the resource for the entire length of the time horizon, their utility will be 4; hence a schedule is only proportionally fair if each agent derives a utility of at least 2 from it. However, any schedule where both agents receive strictly positive utility must have a transfer of the resource, hence only 2 time units are available, and so neither agent can have a utility greater than $\sqrt{2}$.

Note that the schedule allocating each agent one unit of time is Pareto-efficient and envy-free.

Appendix D: Additional Material for Subsection 4.1

D.1. Cadence$_\tau(u)$ is Well-Defined

The next lemma shows that Cadence is well-defined (i.e., the maximizer is finite and unique) for strictly concave functions.

Lemma 3. *Suppose f is increasing, strictly concave, and satisfies $\lim_{t\to\infty} f(t)/t = 0$, then* Cadence$_\tau(f)$ *is well-defined.*

Proof. Cadence$_\tau(f)$ is defined in Eq. (3). Consider the following related optimization problem:

$$\sup_{t \geq 0} F(t), \tag{10}$$

where $F(t) := \frac{f(t)}{t+\tau}$. To show that the cadence is well-defined, we show that the optimal value of Problem (10) is attained, and that the maximizer is unique.

First, we note that F is continuous as the quotient of two continuous functions with domain $\mathbb{R}_{\geq 0}$, where the denominator is strictly positive on $\mathbb{R}_{\geq 0}$ (for all $\tau > 0$). In addition, F is quasiconcave since its upper level sets $S_b^+ := \{t \in \mathbb{R} : F(t) \geq b\} = \{t \in \mathbb{R} : f(t) \geq b(t+\tau)\}$ are convex for all $b \in \mathbb{R}$ by concavity of f. Furthermore, all S_b^+ are in fact strictly convex as intervals in $\mathbb{R}$ (i.e., for any

$t, t' \in S_b^+$ with $t \neq t'$, $\lambda t + (1-\lambda)t'$ lies in the interior of S_b^+ for all $\lambda \in (0,1)$) by strict concavity of f. In addition, the level sets $S_b := \{t \in \mathbb{R} : F(t) = b\} = \{t \in \mathbb{R} : f(t) = b(t+\tau)\}$ have empty interior also by strict concavity of f (in this case, S_b is either empty, a singleton, or the two endpoints of an interval, due to strict concavity of f). It follows that F is strictly quasiconcave.

Now, $F(0) = 0$ and $F(t) \geq 0$ for all $t \geq 0$. By the growth condition $\lim_{t\to\infty} f(t)/t = 0$, for any $\tau > 0$,

$$\lim_{t\to\infty} F(t) = \lim_{t\to\infty} \frac{f(t)\, t}{(t+\tau)\, t} = \left(\lim_{t\to\infty} \frac{t}{t+\tau}\right)\left(\lim_{t\to\infty} \frac{f(t)}{t}\right) = 0,$$

which verifies that F is not unbounded. Since F is continuous, strictly quasiconcave, $F(0) = 0$, and $\lim_{t\to\infty} F(t) = 0$, we conclude that the supremum of F on $\mathbb{R}_{\geq 0}$ is attained and the maximizer is unique. □

D.2. Cadence$_\tau(u)$ for Non-Strictly Concave Utility Functions

Example 2 below shows that if the utility function is not strictly concave, there may not be a finite solution to

$$\arg\max_{t>0} \frac{u(t)}{t+\tau}. \tag{11}$$

Example 2. Suppose $u(t) = t$ (linear utility), then $u(t)/(t+\tau)$ is strictly increasing in t. Hence, the optimal value of Problem (11) is $\sup_{t>0} \frac{u(t)}{t+\tau} = 1$. However, the optimal solution is not attained because there is no finite $t \in \mathbb{R}_{\geq 0}$ that attains this value.

Example 3 below gives a non-strictly concave function where the optimal value is attained, but the optimal solution is not unique.

Example 3. Let

$$u(t) = \begin{cases} t & \text{if } t \in [0,2], \\ \frac{2}{3}(t+1) & \text{if } t \in (2,5], \\ 4 & \text{if } t > 5, \end{cases}$$

and $\tau = 1$. Then, $u(t)/(t+\tau) = 2/3$ for all $t \in [2,3]$, which is the optimal value of $\max_{t>0} \frac{u(t)}{t+\tau}$. It follows that the maximizer is not necessarily unique.

When u is not strictly concave, we can extend the definition of Cadence to be the smallest maximizer of $u(t)/(t+\tau)$, i.e.,

$$\text{Cadence}_\tau(u) = \min\left\{\arg\max_{t>0} \frac{u(t)}{t+\tau}\right\}.$$

D.3. Comparing the Cadence of Two Functions

LEMMA 4. *Let $f:\mathbb{R}\to\mathbb{R}$ and $g:\mathbb{R}\to\mathbb{R}$ be increasing, strictly concave, and differentiable functions, and let $\tau>0$. If the right derivatives of f and g are identical at zero, $f(0)=g(0)$, and $f''(t)\le g''(t)$ for all $0\le t\le \text{CADENCE}_\tau(f)$, then $\text{CADENCE}_\tau(f)\le \text{CADENCE}_\tau(g)$.*

Proof. Define the functions $F(t) := \frac{f(t)}{t+\tau}$ and $G(t) := \frac{g(t)}{t+\tau}$, then $\text{CADENCE}_\tau(f) = \arg\max_{t>0} F(t)$ and $\text{CADENCE}_\tau(g) = \arg\max_{t>0} G(t)$. Since f and g are both strictly concave, $\text{CADENCE}_\tau(f)$ and $\text{CADENCE}_\tau(g)$ are both uniquely defined. The first derivatives of F and G are:

$$F'(t) = \frac{(t+\tau)f'(t) - f(t)}{(t+\tau)^2}$$

and

$$G'(t) = \frac{(t+\tau)g'(t) - g(t)}{(t+\tau)^2},$$

respectively. Since the denominator $(t+\tau)^2$ in both expressions is strictly positive, the first-order optimality condition for $\text{CADENCE}_\tau(f) = \arg\max_{t>0} F(t)$ is equivalent to $\bar{F}'(t) := (t+\tau)f'(t) - f(t) = 0$. Similarly, the first-order optimality condition for $\text{CADENCE}_\tau(g) = \arg\max_{t>0} G(t)$ is equivalent to $\bar{G}'(t) := (t+\tau)g'(t) - g(t) = 0$. Since both $\text{CADENCE}_\tau(f)$ and $\text{CADENCE}_\tau(g)$ are unique, the equations $\bar{F}'(t)=0$ and $\bar{G}'(t)=0$ both have unique solutions. In particular, the functions $\bar{F}'$ and $\bar{G}'$ both only cross zero once and from above.

Next, we have

$$\bar{F}'(0) = \tau f'(0) - f(0) = \tau g'(0) - g(0) = \bar{G}'(0),$$

by assumption that $f(0)=g(0)$ and $f'(0)=g'(0)$, so $\bar{F}'$ and $\bar{G}'$ coincide at $t=0$. Taking the derivatives of $\bar{F}'$ and $\bar{G}'$ gives:

$$\bar{F}''(t) = (t+\tau)f''(t) + f'(t) - f'(t) = (t+\tau)f''(t)$$

and

$$\bar{G}''(t) = (t+\tau)g''(t) + g'(t) - g'(t) = (t+\tau)g''(t).$$

It follows that $\bar{F}''(t)\le \bar{G}''(t)$ for all $t\ge 0$ by assumption that $f''(t)\le g''(t)$ for all $t\ge 0$. Since $\bar{F}'(0)=\bar{G}'(0)$, we must have

$$\bar{F}'(t) = \bar{F}'(0) + \int_0^t \bar{F}''(t)dt \le \bar{G}'(0) + \int_0^t \bar{G}''(t)dt = \bar{G}'(t),\ \forall t\ge 0.$$

In particular, for any $t\ge 0$ with $\bar{F}'(t)=0$ (which satisfies the first-order optimality condition for $\arg\max_{t>0} F(t)$), we have $\bar{G}'(t)\ge 0$. Recall that $\bar{G}'$ only crosses zero once and from above, then if $\bar{G}'(t)=0$ also holds we have $\text{CADENCE}_\tau(f) = \text{CADENCE}_\tau(g)$. Otherwise, if $\bar{G}'(t)>0$ then $\bar{G}'$ has not yet crossed zero and we must have $\text{CADENCE}_\tau(f) < \text{CADENCE}_\tau(g)$. We conclude that $\text{CADENCE}_\tau(f)\le \text{CADENCE}_\tau(g)$. □

D.4. Generalizing Cadence to Discounting and Finite Horizon

The following result shows that $\text{CADENCE}_{(\beta,\tau)}$ and $\text{BOUNDED-CADENCE}_{(T,\tau)}$ are indeed generalizations of CADENCE_{τ}.

LEMMA 5. *Let $\upsilon(t)$ be continuous, strictly positive, and non-increasing. Let $f_\beta(t) := \int_0^t e^{-\beta x}\upsilon(x)dx$ for all $t \geq 0$ for $\beta > 0$, and let $f(t) := \int_0^t \upsilon(x)dx$ for all $t \geq 0$. Then:*
(i) $\lim_{\beta\to 0} \text{CADENCE}_{(\beta,\tau)}(f_\beta) = \text{CADENCE}_{\tau}(f)$.
(ii) $\lim_{T\to\infty} \text{BOUNDED-CADENCE}_{(T,\tau)}(f) = \text{CADENCE}_{\tau}(f)$.

Proof. Let $F(t) := \frac{f(t)}{t+\tau}$. We compute $\text{CADENCE}_{\tau}(f)$ by solving:

$$\max_{t>0} F(t) \equiv \max_{t>0} \frac{f(t)}{t+\tau} \equiv \max_{t>0} \frac{\int_0^t \upsilon(x)dx}{t+\tau}.$$

Since f is strictly concave, F is strictly unimodal and the above optimization problem has a unique optimal solution $t^* = \text{CADENCE}_{\tau}(f)$ by Lemma 3. The function F is differentiable, so t^* is characterized by the first-order optimality condition. The derivative of F is:

$$F'(t) = \frac{\upsilon(t)(t+\tau) - f(t)}{(t+\tau)^2}.$$

Since the denominator of $F'(t)$ is always strictly positive, the condition $F'(t) = 0$ simplifies to $\bar{F}'(t) := \upsilon(t)(t+\tau) - f(t) = 0$.

(i) Let $F_\beta(t) := \frac{f_\beta(t)}{t+\tau}$. We compute $\text{CADENCE}_{(\beta,\tau)}(f_\beta)$ by solving

$$\max_{t>0} F_\beta(t) \equiv \max_{t>0} \frac{f_\beta(t)}{t+\tau} \equiv \max_{t>0} \frac{\int_0^t e^{-\beta x}\upsilon(x)dx}{1-e^{-\beta(t+\tau)}}.$$

Since f_β is strictly concave, F_β is strictly unimodal and the above optimization problem has a unique optimal solution $t^*(\beta) = \text{CADENCE}_{(\beta,\tau)}(f_\beta)$ which depends on β. The functions F_β for all $\beta > 0$ are differentiable, and so $t^*(\beta)$ is characterized by the first-order optimality condition. The derivative of F_β is:

$$F'_\beta(t) = \frac{f'(t)(1-e^{-\beta(t+\tau)}) - \beta e^{-\beta(t+\tau)} f(t)}{(1-e^{-\beta(t+\tau)})^2} = \frac{e^{-\beta t}\left(\upsilon(t)(1-e^{-\beta(t+\tau)}) - \beta e^{-\beta\tau} f(t)\right)}{(1-e^{-\beta(t+\tau)})^2},$$

using $f'(t) = e^{-\beta t}\upsilon(t)$. Since the denominator of $F'_\beta(t)$ is always strictly positive, and $e^{-\beta t}$ is always strictly positive, the condition $F'_\beta(t) = 0$ simplifies to:

$$\bar{F}'_\beta(t) := \upsilon(t)\left(\frac{1-e^{-\beta(t+\tau)}}{\beta}\right) - \int_0^t e^{-\beta(x+\tau)}\upsilon(x)dx = 0.$$

Using l'Hopital's rule gives

$$\lim_{\beta\to 0} \frac{1-e^{-\beta(t+\tau)}}{\beta} = \lim_{\beta\to 0} \frac{(t+\tau)e^{-\beta(t+\tau)}}{1} = t+\tau.$$

Then for all $t \geq 0$ we have

$$\begin{aligned}\lim_{\beta\to 0}\bar{F}'_\beta(t) =& \lim_{\beta\to 0}\left[\upsilon(t)\left(\frac{1-e^{-\beta(t+\tau)}}{\beta}\right)-\int_0^t e^{-\beta(x+\tau)}\upsilon(x)dx\right]\\ =&\upsilon(t)\lim_{\beta\to 0}\left(\frac{1-e^{-\beta(t+\tau)}}{\beta}\right)-\int_0^t\left(\lim_{\beta\to 0}e^{-\beta(x+\tau)}\right)\upsilon(x)dx\\ =&\bar{F}'(t),\end{aligned}$$

where the second equality follows by the dominated convergence theorem, and the third equality uses l'Hopital's rule.

Since t^* is the unique solution to $\bar{F}'(t)=0$, we have $\bar{F}'(t)>0$ for $t<t^*$ and $\bar{F}'(t)<0$ for $t>t^*$. Choose any $\delta>0$, then in particular we have $\bar{F}'(t^*-\delta)>0$ and $\bar{F}'(t^*+\delta)<0$. Let

$$\epsilon = \min\{|\bar{F}'(t^*-\delta)|,|\bar{F}'(t^*+\delta)|\},$$

where $\epsilon>0$. Since $\lim_{\beta\to 0}\bar{F}'_\beta(t^*-\delta)=\bar{F}'(t^*-\delta)$ and $\lim_{\beta\to 0}\bar{F}'_\beta(t^*+\delta)=\bar{F}'(t^*+\delta)$, there is a $\beta'>0$ such that

$$\max\{|\bar{F}'_\beta(t^*-\delta)-\bar{F}'(t^*-\delta)|,|\bar{F}'_\beta(t^*+\delta)-\bar{F}'(t^*+\delta)|\}<\epsilon/2,\ \forall 0<\beta<\beta'.$$

It follows that $\bar{F}'_\beta(t^*-\delta)>\epsilon/2>0$ and $\bar{F}'_\beta(t^*+\delta)<-\epsilon/2<0$. Since $\bar{F}'_\beta(t)=0$ has a unique solution $t^*(\beta)$, $\bar{F}'_\beta$ is continuous, and $\bar{F}'_\beta(t^*-\delta)>0$ and $\bar{F}'_\beta(t^*+\delta)<0$, we must have $t^*(\beta)\in(t^*-\delta,t^*+\delta)$ for all $0<\beta<\beta'$. The choice of δ was arbitrary, so $\lim_{\beta\to 0}t^*(\beta)\to t^*$ must hold, which is the desired result.

(ii) Let $\mathcal{P}_T=(n,f,T,\tau,0)$ for all $T<\infty$. In a tight schedule for $\mathcal{P}_T$ with m intervals all with the same duration, the duration $t>0$ must satisfy $mt+(m-1)\tau=T$ (since there are $m-1$ switches when there are m intervals) or equivalently $m(t+\tau)=T+\tau$. The total utility of this schedule is $mf(t)$, and the corresponding average utility is:

$$\frac{mf(t)}{T}=\frac{T+\tau}{T}\frac{f(t)}{t+\tau},$$

where we divide by T to get the average utility over the time horizon $[0,T]$. Recall $F(t)=\frac{f(t)}{t+\tau}$ and $t^*\in\arg\max_{t>0}F(t)=\textsc{Cadence}_\tau(f)$.

Let

$$\mathcal{X}(T):=\{t>0:(t+\tau)|(T+\tau)\}\equiv\left\{\frac{T-(m-1)\tau}{m}:m=1,\ldots,\lfloor(T+\tau)/\tau\rfloor\right\}$$

be the set of feasible durations for tight schedules of $\mathcal{P}_T$ where all intervals have the same duration (recall that $(t+\tau)|(T+\tau)$ denotes that $t+\tau$ divides $T+\tau$). For each $t\in\mathcal{X}(T)$, the corresponding schedule has average utility $\left(\frac{T+\tau}{t+\tau}\right)\frac{f(t)}{T}=\left(\frac{T+\tau}{T}\right)F(t)$ over $[0,T]$. We maximize the average utility by solving

$$\max_{t\in\mathcal{X}(T)}\left(\frac{T+\tau}{T}\right)F(t)\equiv\left(\frac{T+\tau}{T}\right)\max_{t\in\mathcal{X}(T)}F(t),$$

since we can factor out the constant $\frac{T+\tau}{T}$ in the objective. Then, we see that

$$\textsc{Bounded-Cadence}_{(T,\tau)}(u) = \arg\max_{t \in \mathcal{X}(T)} \left(\frac{T+\tau}{T}\right) F(t) \equiv \arg\max_{t \in \mathcal{X}(T)} F(t),$$

only depends on T through $\mathcal{X}(T)$.

For $T \geq t^*$, we define

$$\begin{aligned} t_l^*(T) &:= \max_m \left\{ \frac{T-(m-1)\tau}{m} \leq t^* : m = 1, \ldots, \lfloor (T+\tau)/\tau \rfloor \right\} \\ &= \frac{T - (\lceil (T+\tau)/(t^*+\tau) \rceil - 1)\tau}{\lceil (T+\tau)/(t^*+\tau) \rceil}, \end{aligned}$$

and

$$\begin{aligned} t_u^*(T) &:= \min_m \left\{ \frac{T-(m-1)\tau}{m} \geq t^* : m = 1, \ldots, \lfloor (T+\tau)/\tau \rfloor \right\} \\ &= \frac{T - (\lfloor (T+\tau)/(t^*+\tau) \rfloor - 1)\tau}{\lfloor (T+\tau)/(t^*+\tau) \rfloor}. \end{aligned}$$

These are the two points in $\mathcal{X}(T)$ that are closest to t^*. Since F is quasi-concave, $\arg\max_{t \in \mathcal{X}(T)} F(t) \in \{t_l^*(T), t_u^*(T)\}$. Note that

$$\lim_{T\to\infty} \frac{T - ((T+\tau)/(t^*+\tau) - 1)\tau}{(T+\tau)/(t^*+\tau)} = \lim_{T\to\infty} \frac{(T+\tau)(t^*+\tau) - (T+\tau)\tau}{(T+\tau)} = t^*,$$

and also that

$$\lim_{T\to\infty} \frac{(T+\tau)/(t^*+\tau)}{\lceil (T+\tau)/(t^*+\tau) \rceil} = 1 \text{ and } \lim_{T\to\infty} \frac{(T+\tau)/(t^*+\tau)}{\lfloor (T+\tau)/(t^*+\tau) \rfloor} = 1.$$

It follows that $\lim_{T\to\infty} t_l^*(T) = t^*$ and $\lim_{T\to\infty} t_u^*(T) = t^*$, which is the desired result.

□

D.5. Computing the Cadence

Lemma 6. *(i)* $\textsc{Cadence}_\tau(f)$ *and* $\textsc{Bounded-Cadence}_{(T,\tau)}(f)$ *can be computed efficiently for non-decreasing and strictly concave* f.

(ii) $\textsc{Cadence}_{(\beta,\tau)}(f)$ *can be computed efficiently for* $f(t) = \int_0^d e^{-\beta x} \upsilon(x) dx$ *where* $\upsilon \geq 0$ *is non-increasing (as in Eq.* (1)*).*

Proof. (i) To compute $\textsc{Cadence}_\tau(f)$, we solve $\max_{t>0} F(t) := f(t)/(t+\tau)$. We see that F is quasi-concave since its upper level sets $S_b^+ = \{t \geq 0 : f(t) \geq b(t+\tau)\}$ are convex for any level $b \geq 0$. It follows that F is unimodal, so we can maximize it with binary search.

$\textsc{Bounded-Cadence}_{(T,\tau)}(f)$ is a maximization problem over the feasible region $\{t : (t+\tau) \mid (T+\tau)\}$. Note that $\textsc{Bounded-Cadence}_{(T,\tau)}(f)$ shares the same objective $F(t)$ as $\textsc{Cadence}_\tau(f)$, which is unimodal. Then, we can similarly do binary search to maximize the unimodal function F over the feasible region $\{t : (t+\tau) \mid (T+\tau)\}$.

(ii) We compute $\textsc{Cadence}_{(\beta,\tau)}(f)$ for $f(t)=\int_0^t e^{-\beta x}\upsilon(x)dx$ for non-increasing υ by solving $\max_{t>0} F(t) := f(t)/(1-e^{-\beta(t+\tau)})$. For threshold $b\geq 0$, the upper level set of F is $S_b^+ = \{t\geq 0 : f(t)\geq b(1-e^{-\beta(t+\tau)})\}$. Let $g_b(t) := f(t) - b(1-e^{-\beta(t+\tau)})$ so that $S_b^+ \equiv \{t\geq 0 : g_b(t)\geq 0\}$. We have $f'(t) = e^{-\beta t}\upsilon(t)$ and

$$g_b'(t) = e^{-\beta t}\upsilon(t) - b\beta e^{-\beta(t+\tau)} = e^{-\beta t}(\upsilon(t) - b\beta e^{-\beta\tau}).$$

We have $e^{-\beta t} > 0$ for all $t\geq 0$ so the sign of $g_b'(t)$ is determined by the expression $\upsilon(t) - b\beta e^{-\beta\tau}$ where $b\beta e^{-\beta\tau} > 0$ is constant in t and $\upsilon(t)\geq 0$ is non-increasing in t. Then the function $g_b'(t)$ only crosses zero from above (i.e., it is first positive and then becomes negative). It follows that $g_b(t)$ is unimodal (it is first non-decreasing and then non-increasing). Consequently, S_b^+ is an interval and is thus a convex set for all $b\geq 0$, so F is quasi-concave. Then, F can be efficiently maximized with binary search. □

D.6. Monotonicity of Cadence

We require the following result for the proof of Proposition 6.

Lemma 7. *Let $\upsilon(t)$ be strictly positive and non-increasing, and let $f_\beta(t) := \int_0^t e^{-\beta x}\upsilon(x)dx$ for all $t\geq 0$ and $\beta > 0$, as in Eq. (1). Then*

(i) $\textsc{Cadence}_\tau(f)$ is strictly increasing in τ.

(ii) $\textsc{Cadence}_{(\beta,\tau)}(f_\beta)$ is strictly increasing in τ, and strictly increasing in β for $\beta\geq 1/\tau$.

Proof. (i) Recall that $\textsc{Cadence}_\tau(f)$ is defined as the solution to the optimization problem:

$$\arg\max_{t>0} F(t) := \frac{f(t)}{t+\tau}.$$

Let $t^*(\tau) = \textsc{Cadence}_\tau(f)$ as a function of τ. The first derivative of the objective F is

$$F'(t) = \frac{(t+\tau)f'(t) - f(t)}{(t+\tau)^2}.$$

Since F is unimodal, the first-order condition $F'(t^*(\tau)) = 0$ is sufficient for optimality. Since the denominator in $F'(t)$ is always strictly positive for $t\geq 0$, the first-order condition simplifies to $(t+\tau)f'(t) = f(t)$. Define $g(t) := \frac{f(t)}{f'(t)}$ and $h(t;\tau) := t+\tau$, then the first-order condition is determined by the intersection $g(t) = h(t;\tau)$.

Note that $g(t)$ and $h(t;\tau)$ for all $\tau\geq 0$ are strictly increasing in t. In particular, g is strictly increasing by the fact that the numerator f is strictly increasing and the denominator f' is non-increasing (by concavity of f). Also notice that $h(0;\tau) = \tau > 0 = g(0)$ for all $\tau > 0$. Since $\textsc{Cadence}_\tau(f)$ is unique by strict concavity of f, $t^*(\tau)$ is the unique solution of $g(t) = h(t;\tau)$. Since $h(0;\tau) > g(0)$, it must be that $g(t) > h(t;\tau)$ for all $t > t^*(\tau)$ and the functions $g(t)$ and $h(t;\tau)$ only cross once. Since $h(t;\tau)$ is strictly increasing in τ for all $t\geq 0$, for any $\tau' > \tau$ we have $g(t^*(\tau)) < h(t^*(\tau);\tau')$. This implies that $t^*(\tau') > t^*(\tau)$, which is the desired result.

(ii) Let $F_\beta(t) := \frac{f_\beta(t)}{1-e^{-\beta(t+\tau)}}$ for all $t \geq 0$, and recall $\text{CADENCE}_{(\beta,\tau)}(f)$ is given by the solution to the optimization problem:

$$\max_{t>0} F_\beta(t) \equiv \max_{t>0} \frac{f_\beta(t)}{1-e^{-\beta(t+\tau)}} \equiv \max_{t>0} \frac{\int_0^t e^{-\beta x}\upsilon(x)dx}{1-e^{-\beta(t+\tau)}}.$$

Let $t^*(\beta,\tau)$ denote the optimal solution to $\text{CADENCE}_{(\beta,\tau)}(f_\beta)$ as a function of $\beta > 0$ and $\tau > 0$. Since log is monotone, we can equivalently maximize $\log(F_\beta(t))$. The first-order condition for the maximization of $\log(F_\beta(\cdot))$ is

$$\frac{d}{dt}\log(F_\beta(t)) = \frac{e^{-\beta t}\upsilon(t)}{\int_0^t e^{-\beta x}\upsilon(x)dx} - \frac{\beta e^{-\beta(t+\tau)}}{1-e^{-\beta(t+\tau)}}.$$

Rearranging terms gives that the equality $\frac{d}{dt}\log(F_\beta(t)) = 0$ holds if and only if:

$$\frac{e^{-\beta t}\upsilon(t)}{\beta e^{-\beta(t+\tau)}} - \frac{\int_0^t e^{-\beta x}\upsilon(x)dx}{1-e^{-\beta(t+\tau)}} = \frac{\upsilon(t)}{\beta e^{-\beta\tau}} - F_\beta(t) = 0.$$

Let $k(\beta,\tau) := \beta e^{-\beta\tau}$, so that $\frac{d}{dt}\log(F_\beta(t)) = 0$ if and only if $\upsilon(t)/k(\beta,\tau) - F_\beta(t) = 0$.

Note that $F_\beta(t)$ is decreasing in β since the numerator is decreasing and the denominator is increasing. Furthermore, $k(\beta,\tau)$ is decreasing in β for $\beta \geq 1/\tau$. Then, $\upsilon(t^*(\beta,\tau))/k(\beta',\tau) - F_{\beta'}(t^*(\beta,\tau))$ is increasing in β' for $\beta \geq 1/\tau$, and it follows that $t^*(\beta,\tau)$ is increasing in β for $\beta \geq 1/\tau$.

Next, note that $k(\beta,\tau)$ is decreasing in τ. Then, $\upsilon(t^*(\beta,\tau))/k(\beta,\tau') - F_\beta(t^*(\beta,\tau))$ is increasing in τ', and it follows that $t^*(\beta,\tau)$ is increasing in τ. □

Appendix E: Additional Material for Subsection 4.2

E.1. $\Phi_{f,\beta,\tau}$ is Well-Defined

LEMMA 8. *Suppose $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is continuous and non-decreasing. Then, $\Phi_{f,\beta,\tau}(t)$ is well-defined for all $t \geq 0$, and $\Phi_{f,\beta,\tau}$ is non-decreasing in t.*

Proof. Fix $t \geq 0$ and consider the equation $f(y) = e^{-\beta(y+\tau)}f(t)$ in the variable $y \geq 0$. If $t = 0$, then $f(t) = f(0) = 0$ and the solution to this equation is $y = 0$. Otherwise, suppose $t > 0$, then the LHS ($f(y)$) is non-decreasing in y by assumption, and the RHS ($e^{-\beta(y+\tau)}f(t)$) is strictly decreasing in y. Since both the LHS and RHS are continuous and $e^{-\beta\tau}f(t) > f(0) = 0$ for $t > 0$, it follows that they cross exactly once. This implies that there is a unique solution to $f(y) = e^{-\beta(y+\tau)}f(t)$ in $y \geq 0$, since otherwise the LHS and RHS would not cross at all or they would cross more than once.

Now, choose $t_1 \leq t_2$ and let $y_1 = \Phi_{f,\beta,\tau}(t_1)$ so that $f(y_1) = e^{-\beta(y_1+\tau)}f(t_1)$. Since f is non-decreasing, we must have

$$f(y_1) = e^{-\beta(y_1+\tau)}f(t_1) \leq e^{-\beta(y_1+\tau)}f(t_2).$$

If $f(t_1) = f(t_2)$, then $f(y_1) = e^{-\beta(y_1+\tau)}f(t_2)$ and $y_1 = \Phi_{f,\beta,\tau}(t_2)$. Otherwise, if $f(t_1) < f(t_2)$ then $f(y_1) < e^{-\beta(y_1+\tau)}f(t_2)$. This implies that $y_1 < \Phi_{f,\beta,\tau}(t_2)$, again using the fact that $e^{-\beta(y+\tau)}f(t)$ is strictly decreasing in y. □

E.2. Comparing Discounted Reciprocal Functions

Lemma 9. *Let $c > 0$ be a real number and let f and g be two non-decreasing functions such that $f(0) = g(0) = 0$. If for every $z \in [0, c]$ it holds that:*

$$\text{for every } t \in [0, z], \frac{g(t)f(z)}{g(z)} \geq f(t),$$

then for any real-valued $\beta, \tau > 0$, $t \in [0, c]$, and any integer $p \geq 1$, $\Phi^p_{f,\beta,\tau}(t) \geq \Phi^p_{g,\beta,\tau}(t)$.

Proof. The proof is by induction. Denote

$$\begin{aligned}\Phi_{f,\beta,\tau}(t) =& \{y \in [0, t] : h_f(y) := f(y)/f(t) = e^{-\beta(y+\tau)}\},\\ \Phi_{g,\beta,\tau}(t) =& \{y \in [0, t] : h_g(y) := g(y)/g(t) = e^{-\beta(y+\tau)}\}.\end{aligned}$$

Note that $\Phi_{f,\beta,\tau}(t)$ and $\Phi_{g,\beta,\tau}(t)$ are always singletons since h_f and h_g are both non-decreasing and $e^{-\beta(y+\tau)}$ is strictly decreasing, so these functions can cross exactly once. Now, for any $t \in [0, c]$, the value of $\Phi_{f,\beta,\tau}(t)$ is given by the intersection between $h_f(y)$ and $e^{-\beta(y+\tau)}$, and the value of $\Phi_{g,\beta,\tau}(t)$ is given by the intersection between $h_g(y)$ and $e^{-\beta(y+\tau)}$ (the function $e^{-\beta(y+\tau)}$ is common to both $\Phi_{f,\beta,\tau}(t)$ and $\Phi_{g,\beta,\tau}(t)$). We can take $z = t$ in the assumption of the lemma to obtain

$$\frac{g(y)f(t)}{g(t)} \geq f(y), \forall y \in [0, t].$$

This immediately implies $h_g(y) \geq h_f(y)$ for all $y \in [0, t]$, and so h_g intersects $e^{-\beta(y+\tau)}$ before h_f does. It follows that $\Phi_{g,\beta,\tau}(t) \leq \Phi_{f,\beta,\tau}(t)$.

Now suppose $\Phi^p_{f,\beta,\tau}(t) \geq \Phi^p_{g,\beta,\tau}(t)$ for all $t \in [0, c]$, we will show it holds for $p + 1$:

$$\begin{aligned}\Phi^{p+1}_{f,\beta,\tau}(t) =& \Phi_{f,\beta,\tau}(\Phi^p_{f,\beta,\tau}(t))\\ \geq& \Phi_{f,\beta,\tau}(\Phi^p_{g,\beta,\tau}(t))\\ \geq& \Phi_{g,\beta,\tau}(\Phi^p_{g,\beta,\tau}(t))\\ =& \Phi^{p+1}_{g,\beta,\tau}(t),\end{aligned}$$

where the first inequality is by monotonicity of $\Phi_{f,\beta,\tau}$ and the induction hypothesis that $\Phi^p_{f,\beta,\tau}(t) \geq \Phi^p_{g,\beta,\tau}(t)$, and the second inequality is by the fact that $\Phi_{f,\beta,\tau}(t) \geq \Phi_{g,\beta,\tau}(t)$. □

Appendix F: Proof of Theorem 1

In this section, we prove Theorem 1. We first give a short outline of the proof. To prove the first bullet point, we require a simple observation, Observation 1, and two preliminary lemmas: Lemmas 10 and 11. Observation 1 states that if all intervals of a schedule have the same length, they can all be allocated to a single agent with no loss of utility. In Lemma 10, we show that if $T = \infty$, there exists some optimal (utilitarian) schedule in which all intervals have the same length.

Taken together with Observation 1, we establish that there is an optimal schedule for $T=\infty$ in which all intervals have the same duration and they are all allocated to the same agent. In Lemma 11, we show that the length of this interval corresponds to the cadence of this agent's utility, by considering a homogeneous problem instance with this utility function. For the second bullet point, we show that when T is as in the statement of Theorem 1, the truncation of the schedule of Lemma 11 at T is optimal. We prove some auxiliary results for the first bullet point in Subsection F.1 and for the second bullet point in Subsection F.2. We complete the proof of Theorem 1 in Subsection F.3.

F.1. Auxiliary Results for the First Bullet Point

The following observation is true because for any $t \geq 0$ and $\mathcal{U}=\{u_1,\ldots,u_n\}$, there exists $i \in [n]$ such that $u_i(t) \geq u_j(t)$ for all $j \in [n]$.

Observation 1. *Let $\mathcal{P}_T=(n,\mathcal{U},T,\tau,\beta)$, for $\beta \geq 0$ and $0 \leq T \leq \infty$, and let π be an arbitrary schedule for $\mathcal{P}_T$ such that $\mathcal{I}(\pi)=\{I_k\}_{k=1}^K$ and $d_k=d_\ell$ for all $k,\ell \in [1,K]$ (where $K=\infty$ for $T=\infty$, and $\{s_k\}_{k\geq 1}$ are arbitrary feasible start times). Then, there exists π' and $i \in [n]$ such that $\pi'_i=\mathcal{I}(\pi)$ and $\pi'_j=\emptyset$ for all $j \neq i$, and*

$$\sum_{i\in[n]} u_i(\pi') \geq \sum_{i\in[n]} u_i(\pi).$$

Lemma 10. *Let $\mathcal{P}_\infty=(n,\mathcal{U},\infty,\tau,\beta)$, where $\beta \geq 0$. There exists a schedule $\pi \in \Pi(\mathcal{P})$ that is optimal for $\mathcal{P}_\infty^{UC}$ where $\mathcal{I}(\pi)=\{I_k\}_{k\geq 1}$, and $I_k=(s_k,d)$ for some $d>0$ for all $k \geq 1$.*

Proof. Towards a contradiction, suppose there is no such optimal schedule (and so all optimal schedules have at least two intervals with different durations). Let π be an optimal schedule for $\mathcal{P}_\infty^{UC}$, set $\mathcal{I}=\{I_k\}_{k\geq 1}$ where $I_k=(s_k,d_k)$ for all $k \geq 1$, and let ν be the smallest integer such that: $d_\nu \neq d_{\nu+1}$ and $d_j=d_{j+1}$ for all $j \in \{1,\ldots,\nu-1\}$. Let $A=\{I_1,I_2,\ldots,I_{\nu-1}\}$ and $B=\{I_{\nu+1},I_{\nu+2},\ldots\}$; in other words, $\mathcal{I}(\pi)=A\cup\{I_\nu\}\cup B$. Denote $d^*=d_1=\cdots=d_\nu$.

By Observation 1, we can assume w.l.o.g. that $A \subset \pi_1$ (i.e., that all intervals in A are allocated to agent 1). Then $u_1(A)=\sum_{j=1}^{\nu-1} e^{-\beta(j-1)(d^*+\tau)}u_1(d^*)$ is the total utility that agent 1 gets from A, and $u_1(I_\nu)=e^{-\beta(\nu-1)(d^*+\tau)}u_1(d^*)$ is the utility that agent 1 gets from I_ν. For all $j \geq \nu+1$, let $\zeta(j)$ denote the agent to which I_j is assigned in π. Let $B^\dagger=\{I_k^\dagger\}_{k\geq \nu+1}$, $I_k^\dagger=(s_k-s_{\nu+1},d_k)$ be the set of intervals identical to B except that their start times have been reduced by $s_{\nu+1}$ (in other words, $B^\dagger$ is the set of intervals of $\mathcal{I}$ if time started at $s_{\nu+1}$ instead of 0). Let $u(B^\dagger)$ denote the total utility that would be derived from $B^\dagger$ if the intervals were allocated to the same agents as in π (i.e., if for each $j \geq \nu+1$, $I_j^\dagger$ is allocated to $\zeta(j)$).

Let $\mathcal{I}'=\{I'_1,I'_2,\ldots,I'_{\nu-1},I'_{\nu+1},I'_{\nu+2},\ldots\}$ be such that for all $j<\nu$, $I'_j=I_j$, and for all $j \geq \nu+1$, $I'_j=(s_j-d^*-\tau,d_j)$. Simply put, $\mathcal{I}'$ is obtained from $\mathcal{I}$ by removing I_ν and shifting all of the intervals after I_ν to the left, so that the schedule is tight. Let $\pi' \in \Pi(\mathcal{P})$ be a schedule for $\mathcal{I}'$ where each interval I'_j is allocated to the same agent as I_j is allocated in π. Then

$$u(\pi) = u_1(A) + u_1(I_\nu) + e^{-\beta\nu(d^*+\tau)}u(B^\dagger) > u(A) + e^{-\beta(\nu-1)(d^*+\tau)}u(B^\dagger) = u(\pi'), \tag{12}$$

otherwise π' would be an optimal schedule for which the durations of the $(\nu-1)^{th}$ and ν^{th} intervals are different, which would contradict the choice of π. Simplifying Eq. (12), we get

$$u_1(d^*) + e^{-\beta(d^*+\tau)}u(B^\dagger) > u(B^\dagger). \tag{13}$$

Next consider the schedule $\pi^\ddagger$ where $\mathcal{I}(\pi^\ddagger) = A \cup \{I_\nu, I_\nu^\ddagger\} \cup B^\ddagger$, $I_\nu^\ddagger = (s_\nu + d^* + \tau, d^*)$, and $B^\ddagger = \{I_{\nu+1}^\ddagger, I_{\nu+2}^\ddagger, \ldots\}$ where $I_i^\ddagger = (s_i + d^* + \tau, d_i)$ for $i \geq \nu+1$, and $I_i^\ddagger$ is allocated to the same agent as I_i is under π. The set $B^\ddagger$ is identical to B, except that the intervals are shifted to the right by $d^*+\tau$. $B^\ddagger$ is also identical to $B^\dagger$, except the intervals are all shifted to the left by $(\nu+1)d^*+\tau$. Then,

$$\begin{aligned} u(\pi^\ddagger) &= u_1(A) + u_1(I_\nu) + e^{-\beta(d^*+\tau)}u(I_\nu) + e^{-\beta(\nu+1)(d^*+\tau)}u(B^\dagger) \\ &> u_1(A) + u_1(I_\nu) + e^{-\beta\nu(d^*+\tau)}u(B^\dagger) = u(\pi), \end{aligned}$$

where the inequality is due to Eq. (13). Therefore $u(\pi^\ddagger) > u(\pi)$, which contradicts the optimality of π. □

For $i \in [n]$, define $\mathcal{P}_{\infty,i}^{UC} = (n, u_i, \infty, \tau, \beta)$. We show that for any $i \in [n]$, there is an optimal utilitarian schedule for $\mathcal{P}_{\infty,i}^{UC}$ where all intervals have the same duration, and that this optimal duration is precisely the cadence of u_i.

Lemma 11. *Let $\mathcal{P}_{\infty,u} = (n, u, \infty, \tau, \beta)$, $n \in \mathbb{N}_{\geq 1}$, $\tau > 0$, $\beta > 0$, and $\delta = \text{Cadence}_{(\beta,\tau)}(u)$. Set $I_k = (s_k, d_k)$ with $s_k = (k-1)(\delta+\tau)$ and $d_k = \delta$ for all $k \geq 1$. Then any $\pi = \{\pi_1, \ldots, \pi_n\}$ with $\mathcal{I}(\pi) = \{I_k\}_{k\geq 1}$ is an optimal schedule for $\mathcal{P}_{\infty,u}^{UC}$. Furthermore, $OPT\left(\mathcal{P}_{\infty,u}^{UC}\right) = \frac{u(\delta)}{1-e^{-\beta(\delta+\tau)}}$.*

Proof. By Lemma 10, there is an optimal schedule π where all intervals have the same duration. We show that the schedule is optimal if the common duration is $\delta = \text{Cadence}_{(\beta,\tau)}(u)$. Given $OPT\left(\mathcal{P}_{\infty,u}^{UC}\right)$, let δ^* be an optimal solution of

$$\max_{t_1\geq 0}\left\{u(t_1) + e^{-\beta(t_1+\tau)}OPT\left(\mathcal{P}_{\infty,u}^{UC}\right)\right\}.$$

Then $OPT\left(\mathcal{P}_{\infty,u}^{UC}\right) = u(\delta^*) + e^{-\beta(\delta^*+\tau)}OPT\left(\mathcal{P}_{\infty,u}^{UC}\right)$, or equivalently

$$OPT\left(\mathcal{P}_{\infty,u}^{UC}\right) = \frac{u(\delta^*)}{1-e^{-\beta(\delta^*+\tau)}}.$$

For $\delta = \text{Cadence}_{(\beta,\tau)}(u)$,

$$\frac{u(\delta)}{1-e^{-\beta(\delta+\tau)}} = \max_{t\geq 0}\frac{u(t)}{1-e^{-\beta(t+\tau)}}.$$

By the definition of δ,

$$OPT\left(\mathcal{P}_{\infty,u}^{UC}\right) = \frac{u(\delta^*)}{1-e^{-\beta(\delta^*+\tau)}} \leq \frac{u(\delta)}{1-e^{-\beta(\delta+\tau)}}. \tag{14}$$

Since $\frac{u(\delta)}{1-e^{-\beta(\delta+\tau)}}$ is the total discounted utility of the schedule where all intervals have duration δ,

$$\frac{u(\delta)}{1-e^{-\beta(\delta+\tau)}} \leq OPT\left(\mathcal{P}^{UC}_{\infty,u}\right). \tag{15}$$

Therefore, from Ineq. (14) and (15),

$$OPT\left(\mathcal{P}^{UC}_{\infty,u}\right) = \frac{u(\delta)}{1-e^{-\beta(\delta+\tau)}},$$

and so the schedule where all intervals have duration $\delta = \text{CADENCE}_{(\beta,\tau)}(u)$ is optimal. □

F.2. Auxiliary Results for the Second Bullet Point

The following lemma describes the loss of utility due to truncating the optimal utilitarian schedule for an infinite horizon at some finite T such that $T = m(\delta_{i^*}+\tau)-\tau$, where m is a (sufficiently large) integer.

LEMMA 12. *Let* $\mathcal{P}_\infty = (n,\mathcal{U},\infty,\tau,\beta)$, $n \in \mathbb{N}_{\geq 1}$, $\tau > 0$, $\beta > 0$, $\delta_i = \text{CADENCE}_{(\beta,\tau)}(u_i)$, $i^* \in \arg\max_{i\in[n]} \frac{u_i(\delta_i)}{1-e^{-\beta(\delta_i+\tau)}}$, *and* $\rho = e^{-\beta(\delta_{i^*}+\tau)}$. *Let* $m > n$ *be some integer, set* $T = m(\delta_{i^*}+\tau)-\tau$, *and let* $\mathcal{P}_T = (n,\mathcal{U},T,\tau,\beta)$. *Then*

$$OPT\left(\mathcal{P}^{UC}_{T}\right) = (1-\rho^m)OPT\left(\mathcal{P}^{UC}_{\infty}\right). \tag{16}$$

Proof. To demonstrate equality, we show that each side of Eq. (16) is greater than or equal to the other side.

- $OPT\left(\mathcal{P}^{UC}_{T}\right) \geq (1-\rho^m)OPT\left(\mathcal{P}^{UC}_{\infty}\right)$: Let π be an optimal schedule for $\mathcal{P}^{UC}_{\infty}$, in which each interval in $\mathcal{I}(\pi)$ has duration δ_{i^*}; one such schedule must exist by Lemma 11. For any $T < \infty$, let $\pi^T = \{\pi_i^T\}_{i\in[n]}$ be the truncation of π to include only those intervals that end at or before time T, i.e., where $\pi_i^T := \{I_k \in \pi_i : s_k + d_k \leq T\}$ for all $i \in [n]$.

  Then we have

$$\begin{aligned} OPT\left(\mathcal{P}^{UC}_{T}\right) &\geq u_{i^*}(\pi^T,\mathcal{P}_T) \\ &= \sum_{k=0}^{m-1} \rho^k u_{i^*}(\delta_{i^*}) \\ &= u_{i^*}(\delta_{i^*})\frac{1-\rho^m}{1-\rho}. \end{aligned} \tag{17}$$

  From Lemma 11, $OPT\left(\mathcal{P}^{UC}_{\infty}\right) = \frac{u_{i^*}(\delta_{i^*})}{1-\rho}$. Combining this equality with Ineq. (17) gives the inequality.
- $OPT\left(\mathcal{P}^{UC}_{T}\right) \leq (1-\rho^m)OPT\left(\mathcal{P}^{UC}_{\infty}\right)$: Assume the contrary, i.e., $OPT\left(\mathcal{P}^{UC}_{T}\right) > (1-\rho^m)OPT\left(\mathcal{P}^{UC}_{\infty}\right)$. Let π be an optimal schedule for $\mathcal{P}^{UC}_{\infty}$ with $\mathcal{I}(\pi) = \{I_1, I_2, \ldots\}$ and let π^T be an optimal schedule for $\mathcal{P}^{UC}_{T}$. For integer m as in the statement of the lemma, let $\pi^{T:\infty}$ be the schedule π restricted to the tail intervals $\{I_{m+1}, I_{m+2}, \ldots\}$. Then, define the new schedule $\pi^\dagger = \pi^T \cup \pi^{T:\infty}$ which is feasible for $\mathcal{P}^{UC}_{\infty}$. However, $\sum_{i\in[n]} u_i(\pi^\dagger) > \sum_{i\in[n]} u_i(\pi)$, contradicting the optimality of π for $\mathcal{P}^{UC}_{\infty}$. □

F.3. Putting It All Together

THEOREM 1. *Let $\mathcal{P} = (n, \mathcal{U}, T, \tau, \beta)$, $n \in \mathbb{N}_{\geq 1}$, $\tau > 0$, $\beta > 0$, and for each $i \in [n]$, denote $\delta_i = \text{CADENCE}_{(\beta,\tau)}(u_i)$ and $\rho_i = e^{-\beta(\delta_i+\tau)}$. Set $i^* \in \arg\max_{i\in[n]} \frac{u_i(\delta_i)}{1-\rho_i}$. Let $I_k = (s_k, \delta_{i^*})$ with $s_k = (k-1)(\delta_{i^*}+\tau)$ for all k such that $s_k + \delta_{i^*} \leq T$. Set $\pi = \{\pi_1, \ldots, \pi_n\}$, where $\pi_{i^*} = \{I_k\}_{k\geq 1}$ and $\pi_j = \emptyset$ for all $j \neq i^*$. For any $m \in \mathbb{Z}$, denote $T_m = m(\delta_{i^*}+\tau) - \tau$.*

- *If $T = T_m$ for some $m \in \mathbb{Z}$, then π is an optimal schedule for $\mathcal{P}^{UC}$.*
- *Otherwise, for any $\epsilon > 0$, there exists $m_\epsilon \in \mathbb{Z}_{>0}$ such that if $T > T_{m_\epsilon}$ then*

$$\sum_{i\in[n]} u(\pi) \geq (1-\epsilon) OPT\left(\mathcal{P}^{UC}\right).$$

Proof. We prove each bullet point separately.

First bullet point: Lemma 11 characterizes an optimal schedule for homogeneous agents when $T = \infty$. Consider now the n optimal schedules for $\mathcal{P}_{\infty,i}$, for $i \in [n]$, where $\mathcal{P}_{\infty,i} = (n, u_i, T, \tau, \beta)$. From Lemma 10 and Observation 1, we know that there is some agent i for which the schedule for $\mathcal{P}_{\infty,i}$ is also optimal for $\mathcal{P}_\infty$. From Lemma 11, we know that $OPT\left(\mathcal{P}^{UC}_{\infty,i}\right) = \frac{u(\delta_i)}{1-e^{-\beta(\delta_i+\tau)}}$, and so clearly, from the definition of i^* in the statement of Theorem 1, the optimal schedule for $\mathcal{P}_{\infty,i^*}$ is also optimal for $\mathcal{P}_\infty$. To complete the proof, we show that when $T = T_m$, the truncation of the schedule of Lemma 11 for $\mathcal{P}_{\infty,i^*}$ at T is optimal. Formally, set $u = u_{i^*}$, and let $\mathcal{P}_u = (n, u, T, \tau, \beta)$ and $\mathcal{P}_{\infty,u} = (n, u, \infty, \tau, \beta)$. For all $k \geq 1$, let $I_k = (s_k, \delta)$ with $s_k = (k-1)(\delta+\tau)$, where $\delta = \text{CADENCE}_{(\beta,\tau)}(u)$. Let π^∞ be an optimal schedule for $\mathcal{P}^{UC}_{\infty,u}$; that is, $\mathcal{I}(\pi^\infty) = \{I_k\}_{k\geq 1}$. Let m be as in the statement of Theorem 1; that is, $T = m(\delta+\tau) - \tau$ for some integer m. Let π^T be such that $\mathcal{I}(\pi^T) = \{I_1, I_2, \ldots, I_m\}$, suppose for a contradiction that π^T is not optimal for $\mathcal{P}^{UC}_u$, and denote the optimal schedule for $\mathcal{P}^{UC}_u$ by π^*.

Let $\pi^{T:\infty}$ be the schedule π^∞ restricted to the tail intervals $\{I_{m+1}, I_{m+2}, \ldots\}$. Define a new schedule $\pi^\dagger = \pi^* \cup \pi^{T:\infty}$. This schedule is feasible for $\mathcal{P}^{UC}_{\infty,u}$. However, $\sum_{i\in[n]} u_i(\pi^\dagger) > \sum_{i\in[n]} u_i(\pi)$, contradicting the optimality of π^∞ for $\mathcal{P}^{UC}_{\infty,u}$.

Second bullet point: For any $\epsilon > 0$, set $m_\epsilon = \frac{\ln(\epsilon)}{\ln(\rho_{i^*})}$. For any integer $m > 0$, define $T_m = m(\delta_{i^*}+\tau) - \tau$, and set $\mathcal{P}_m = (n, u, T_m, \tau, \beta)$. From Lemma 12, $OPT\left(\mathcal{P}^{UC}_{m_\epsilon}\right) = (1-\rho^{m_\epsilon}) OPT\left(\mathcal{P}^{UC}_\infty\right)$. As $OPT\left(\mathcal{P}^{UC}_\infty\right) \geq OPT\left(\mathcal{P}^{UC}\right)$,

$$\sum_{i\in[n]} u(\pi) \geq OPT\left(\mathcal{P}^{UC}_{m_\epsilon}\right) = (1-\epsilon) OPT\left(\mathcal{P}^{UC}_\infty\right) \geq (1-\epsilon) OPT\left(\mathcal{P}^{UC}\right),$$

where π is the schedule defined in the theorem statement. □

The next corollary to Theorem 1 follows immediately from the above proof and Lemma 10.

COROLLARY 3. *Let $\mathcal{P}_\infty = (n, \mathcal{U}, \infty, \tau, \beta)$, $n \in \mathbb{N}_{\geq 1}$, $\tau > 0$, $\beta \geq 0$, and let $\mathcal{P}_{\infty,i} = (n, u_i, \infty, \tau, \beta)$ for each $i \in [n]$. Then $OPT\left(\mathcal{P}^{UC}_\infty\right) = \max_{i\in[n]} OPT\left(\mathcal{P}^{UC}_{\infty,i}\right)$.*

Appendix G: Pseudocode for the Greedy Algorithm

Algorithm 2: GREEDY

```
Input: P = (n, u, T, τ, β)
Output: A schedule π.
δ ← CADENCE_(β,τ)(u)
ρ ← e^{-β(δ+τ)}
m = max{k : k(δ+τ) + δ ≤ T}
for i ∈ [n] do
  | π_i = ∅
end
for k ∈ [0, m] do
  | j ← arg min_{i∈[n]} u_i(π)
  | π_j ⇐ π_j ∪ (k(δ+τ)}, δ)
end
return π = {π_1, ..., π_n}
```

Appendix H: Proof of Theorem 2

In this section, we prove Theorem 2. We first give a short outline of the proof. To prove the first bullet point of Theorem 2, we introduce a new definition, (m,n)-envy-freeness: For problem $\mathcal{P}$, a schedule π is (m,n)-*envy-free* for $m \leq n$ if $u_i(\pi_i) \geq u_i(\pi_j)$ for all $i \in [m], j \in [n]$. In other words, a schedule is (m,n)-envy-free if at least m agents do not prefer another agent's allocation to their own. In the homogeneous setting, this definition implies that $u(\pi_i) = u(\pi_j)$ for all $i,j \in [m]$.

We consider a family of problem instances that is auxiliary to $\mathcal{P}$, (where $\mathcal{P} = (n,u,\infty,\tau,\beta)$ is as in the theorem statement). For each $m \leq n$, let $\mathcal{P}_m = (m,u,\infty,\tau,\beta)$ be identical to $\mathcal{P}$, except that it only contains the agents $1,\dots,m$. Let $\mathcal{P}^{EF}_{(m,n)}$ denote the following problem: compute a schedule $\pi \in \Pi(\mathcal{P})$ that maximizes the sum of (all of) the agents' utilities, subject to π being (m,n)-envy-free. The advantage of this definition is that it allows us to use induction, as it does not impose the assumption that all agents have the same utility in an optimal schedule. We show in Lemma 14 (below) that the output of Algorithm 1, when it is executed on $\mathcal{P}_m$, is an optimal (m,n)-envy-free schedule. Setting $m = n$ completes the proof.

We first prove a technical lemma.

LEMMA 13. *Let $\beta \in \mathbb{R}_{>0}$, and let $u : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be a utility function in the form of Eq.* (1) *where $u(t) = \int_0^t e^{-\beta x} \upsilon(x) dx$ for all $t \geq 0$. For any constant $y \in \mathbb{R}_{>0}$, let $f(t,y) := u(t) + y e^{-\beta(t+\tau)}$.*

(i) The function $f(t,y)$ is unimodal in t.

(ii) The function $f(t,y)$ attains its maximum at $t^*(y)$ where $\upsilon(t^*(y)) = \beta y e^{-\beta\tau}$, and $t^*(y)$ is non-increasing in y.

Proof. (i) Since $u(t) = \int_0^t e^{-\beta x} \upsilon(x) dx$ for all $t \geq 0$, we have $u'(t) = e^{-\beta t}\upsilon(t)$. Differentiate f w.r.t. t to obtain:

$$\frac{d}{dt}\left[u(t) + y e^{-\beta(t+\tau)}\right] = e^{-\beta t}\upsilon(t) - \beta y e^{-\beta(t+\tau)} = e^{-\beta t}\left(\upsilon(t) - \beta y e^{-\beta\tau}\right).$$

This expression is the product of the function $e^{-\beta t}$, which is strictly positive on $t > 0$, and the function $\upsilon(t) - \beta y e^{-\beta\tau}$. Note that $\beta y e^{-\beta\tau}$ is a constant and υ is non-increasing. Then, $\upsilon(t) - \beta y e^{-\beta\tau}$ only crosses zero at most once and from above (it is first positive and then negative). It follows that f is first non-decreasing and then non-increasing, and so it is unimodal.

(ii) By the previous part, the necessary condition for f to reach its maximum in t for fixed y is $\frac{\partial}{\partial t} f(t,y) = 0$. Then, we must have $\upsilon(t) = \beta y e^{-\beta\tau}$ at any stationary point t of f, which must also be a maximizer of f since it is unimodal. Since $\upsilon(t)$ is non-increasing, the solution $t^*(y)$ is non-increasing in y. □

The following lemma is the main building block in our proof that Algorithm 1 computes an optimal fair schedule.

LEMMA 14. *For any $\mathcal{P} = (n, u, \infty, \tau, \beta)$ the output of Algorithm 1 on $\mathcal{P}_m = (m, u, \infty, \tau, \beta)$ is optimal for $\mathcal{P}^{EF}_{(m,n)}$.*

Proof. The proof is by induction on m. Note that the values of δ, ρ, and λ^* that are computed in Algorithm 1 are independent of m ($\delta = \text{CADENCE}_{(\beta,\tau)}(u)$, $\rho = e^{-\beta(\delta+\tau)}$, and $\lambda^* = \max_{\lambda\in[n]} \left\{\lambda : \rho \geq \frac{\lambda-1}{\lambda}\right\}$). The base cases are the set of values $\{m : m \leq \lambda^*\}$ where we have $\rho > \frac{m-1}{m}$. Then, by Corollary 2, the greedy allocation of the intervals of the optimal schedule of $\mathcal{P}^{UC}$ generates an envy-free schedule, as the corollary implies a partition of the intervals into sets $\pi_1, \ldots, \pi_n$ such that $u(\pi_i) = \frac{u_i(\delta)}{m(1-\rho)}$ for each $i \in [m]$, and $\pi_{m+1}, \ldots, \pi_n$ are empty sets. Note that in this case, Algorithm 1 simply returns the schedule generated by Algorithm 2 (i.e., it does not execute any push right operation and it does not truncate the schedule). As the total utility of this schedule is equal to the total utility of the optimal schedule and it is envy-free for the first m agents, this schedule is optimal for $\mathcal{P}^{EF}_{(m,n)}$.

For the inductive step, we assume that the lemma statement holds for $m-1$ agents; i.e., that the output of Algorithm 1 on $\mathcal{P}_{m-1} = (m-1, u, \infty, \tau, \beta)$ is optimal for $\mathcal{P}^{EF}_{(m-1,n)}$. We show that it holds for m agents as well. For any schedule π, let $u_{[k]}(\pi) = \sum_{i\in[k]} u_i(\pi)$ denote the total utility of agents $1, \ldots, k$ under π. Let π_{m-1} and π_m denote the schedules output by Algorithm 1 for $\mathcal{P}_{m-1}$ and $\mathcal{P}_m$, respectively. By the induction hypothesis, π_{m-1} is optimal for $\mathcal{P}^{EF}_{(m-1,n)}$. Let π^*_m denote the optimal fair schedule for $\mathcal{P}^{EF}_{(m,n)}$, and let δ_m and δ^*_m denote the duration of the first interval in π_m and π^*_m, respectively. When $T = \infty$ and $m > \lambda^*$, Algorithm 1 executed on m agents allocates the first interval to agent m, and the remaining $m-1$ agents are allocated π_{m-1}, pushed right by $\delta_m + \tau$. If the first interval is identical in both π_m and π^*_m, then the schedules must be identical (otherwise π^*_m allocates more utility than the utility of π_{m-1} for the remaining intervals, contradicting the inductive hypothesis). Therefore, in order to verify that $\pi_m = \pi^*_m$, it suffices to show that:

1. $\delta^*_m \leq \delta_m$,

2. $\delta_m^* \geq \delta_m$.

Intuitively, if the first inequality does not hold then, as π_m^* allocates a larger interval to agent m, there is simply not enough utility available to be allocated to the remaining agents, hence the resulting schedule is not envy-free. Establishing the second inequality is slightly more involved; we show that the total utility of the schedule is unimodal in the length of the first interval, and that it is increasing at δ_m^*.

Towards a contradiction to the first inequality, assume that $\delta_m^* > \delta_m$. Then, assuming w.l.o.g. that agent m is allocated the first interval in both schedules, the utility of the first $m-1$ agents in π_m^* can be bounded as follows:

$$\begin{aligned} u_{[m-1]}(\pi_m^*) &\leq e^{-\beta(\delta_m^*+\tau)} u_{[m-1]}(\pi_{m-1}) \\ &< e^{-\beta(\delta_m+\tau)} u_{[m-1]}(\pi_{m-1}) \\ &= u_{[m-1]}(\pi_m). \end{aligned} \tag{18}$$

The first inequality is because π_{m-1} is an optimal schedule in which agents $1, \dots, m-1$ receive all of the intervals, the second is due to the assumption that $\delta_m^* > \delta_m$.

As both π_m and π_m^* are (m,n)-envy free, it holds that $u_1(\pi_m^*) = \frac{u_{[m-1]}(\pi_m^*)}{m-1}$ and $u_1(\pi_m) = \frac{u_{[m-1]}(\pi_m)}{m-1}$, but then, combining with Eq. (18), we have that $u_{[m]}(\pi_m) > u_{[m]}(\pi_m^*)$, contradicting the optimality of π_m^*.

For the second inequality, assume that $\delta_m^* < \delta_m$. We again assume w.l.o.g. that the first interval is allocated to agent m in both schedules, and note that it is possible that agent m is allocated additional intervals.

Define $f(t,y) = u(t) + e^{-\beta(t+\tau)} y$. By Algorithm 1,

$$u_{[n]}(\pi_m) = u(\delta_m) + e^{-\beta(\delta_m+\tau)} u_{[n]}(\pi_{m-1}) = f\left(\delta_m, u_{[n]}(\pi_{m-1})\right). \tag{19}$$

As the length of the first interval in π_m^* is δ_m^*, the optimal total utility of all n agents (i.e., including agent m) for all remaining intervals is upper bounded by

$$e^{-\beta(\delta_m^*+\tau)} u_{[m-1]}(\pi_{m-1}).$$

This is because $u_{[m-1]}(\pi_{m-1})$ is the optimal utility for a fair schedule when $m-1$ agents are allocated equal shares. We note that this is where we require (m,n)-envy-freeness, as we do not exclude the possibility here that other agents, such as agent m, are also allocated some intervals in π_{m-1}. Therefore,

$$\begin{aligned} u_{[n]}(\pi_m^*) &\leq u(\delta_m^*) + e^{-\beta(\delta_m^*+\tau)} u_{[m-1]}(\pi_{m-1}) \\ &= f\left(\delta_m^*, u_{[n]}(\pi_{m-1})\right). \end{aligned} \tag{20}$$

Eq. (20) holds because $u_{[m-1]}(\pi_{m-1}) = u_{[n]}(\pi_{m-1})$ as the agents $m, m+1, \ldots, n$ are not allocated any intervals in π_{m-1} under Algorithm 1.

Let $t^* := \arg\max_{t>0} f(t, u_{[n]}(\pi_{m-1}))$. By Lemma 13, f is unimodal. Therefore, if $\delta_m^* < \delta_m \leq t^*$, then $f(\delta_m^*, u_{[n]}(\pi_{m-1})) < f(\delta_m, u_{[n]}(\pi_{m-1}))$. Taken with Eq. (19) and Ineq. (20), this will imply that $u_{[n]}(\pi_m) > u_{[n]}(\pi_m^*)$, in contradiction to the optimality of π_m^*.

To establish that $\delta_m \leq t^*$, we show that $\delta_m \leq \textsc{Cadence}_{(\beta,\tau)}(u) \leq t^*$. From Lemma 11, we know that the function $f(t, OPT(\mathcal{P}^{UC}))$ attains it maximum at $\textsc{Cadence}_{(\beta,\tau)}(u)$, as the optimal schedule for $\mathcal{P}^{UC}$ has all intervals of length $\textsc{Cadence}_{(\beta,\tau)}(u)$. As $u_{[n]}(\pi_m) \leq OPT(\mathcal{P}^{UC})$ (because $u_{[n]}(\pi_m)$ is upper bounded by the optimal total utility $OPT(\mathcal{P}^{UC})$), we have that $t^* \geq \textsc{Cadence}_{(\beta,\tau)}(u)$ by Lemma 13 part (ii). Finally, as $\rho < \frac{m-1}{m}$, it must hold that $\delta_m \leq \textsc{Cadence}_{(\beta,\tau)}(u)$, as otherwise the total utility allocated to the remaining $m-1$ agents would be strictly less than $(m-1)u(\delta_m)$. □

H.1. Putting it all together

THEOREM 2. *Let $\mathcal{P} = (n, u, T, \tau, \beta)$, $n \in \mathbb{N}_{\geq 1}$, $\tau > 0$, $\beta > 0$, $\delta = \textsc{Cadence}_{(\beta,\tau)}(u)$, and $\rho = e^{-\beta(\delta+\tau)}$. Let π be the schedule generated by Algorithm 1. Then π is envy-free, and:*

- *If $T = \infty$ then π is an optimal envy-free schedule for $\mathcal{P}$.*
- *Otherwise, for any $\epsilon > 0$, there exists $T_\epsilon > 0$ such that if $T > T_\epsilon$ then*

$$\sum_{i \in [n]} u(\pi) \geq (1-\epsilon) OPT\left(\mathcal{P}^{EF}\right).$$

Proof. We prove each bullet point separately.

First bullet point: By setting $m = n$ in Lemma 14 gives that the output of Algorithm 1 on $\mathcal{P} = (m, u, \infty, \tau, \beta)$ is optimal for $\mathcal{P}^{EF}_{(n,n)}$, and hence is the optimal envy-free schedule.

Second bullet point: Let $\pi^\dagger$ denote the schedule after the first loop in Algorithm 1; that is, $\pi^\dagger$ is the schedule before the truncation at T. Let $\pi^\dagger_{[\lambda^*]}$ be $\pi^\dagger$, restricted to the first λ^* agents, that is, $\pi^\dagger_{[\lambda^*]} = \{\pi^\dagger_1, \ldots, \pi^\dagger_{\lambda^*}\}$. Let t_{λ^*} be the start time of the first interval in $\pi^\dagger_{[\lambda^*]}$. Let $\mathcal{P}_{t_{\lambda^*},\infty} = (n, u, [t_{\lambda^*}, \infty), \tau, \beta)$ denote the problem that is identical to $\mathcal{P}$ except the time horizon is $[t_{\lambda^*}, \infty)$ instead of $[0, T]$. It is straightforward to adapt the proof of Theorem 1 to see that $\pi^\dagger_{[\lambda^*]}$ is an optimal utilitarian schedule for $\mathcal{P}_{t_{\lambda^*},\infty}$, and similarly, that there exists some T_ϵ such that if $\pi^\dagger_{[\lambda^*]}$ is truncated at T_ϵ (denote this truncated schedule by $\pi^\ddagger_{[\lambda^*]}$) then

$$\sum_{i \in [\lambda^*]} u_i(\pi^\ddagger_{[\lambda^*]}) > \left(1 - \frac{\epsilon}{n}\right) u_i(\pi^\dagger_{[\lambda^*]}).$$

By the first bullet point, $\pi^\dagger$ is the optimal schedule for $\mathcal{P}^{EF}_\infty$, where $\mathcal{P}_\infty = (n, u, \infty, \tau, \beta)$. As the utility of each agent is reduced by at most $\frac{\epsilon}{n}$ after truncation of $\pi^\dagger$, we have that

$$\sum_{i \in [n]} u(\pi) \geq (1-\epsilon) OPT(\mathcal{P}^{EF}_\infty) \geq (1-\epsilon) OPT\left(\mathcal{P}^{EF}\right),$$

where π is the schedule defined in the theorem statement. □

Appendix I: Proof of Theorem 3 and Corollary 1

In order to prove Theorem 3, we would like to compare $OPT\left(\mathcal{P}^{UC}\right)$ and $OPT\left(\mathcal{P}^{EF}\right)$. However, it is unclear how to directly compare these two values. We therefore use several auxiliary problems and a chain of their relationships. For any problem $\mathcal{P}=(n,\mathcal{U},T,\tau,\beta)$, we denote $\mathcal{P}_T=\mathcal{P}$ for clarity and define the following variants. Let $\mathcal{P}_\infty=(n,\mathcal{U},\infty,\tau,\beta)$ be the problem with the same parameters as $\mathcal{P}$, except that the time horizon is infinite. For each $i\in[n]$, let $\mathcal{P}_{T,i}=(n,u_i,T,\tau,\beta)$ and $\mathcal{P}_{\infty,i}=(n,u_i,\infty,\tau,\beta)$ be the variants of $\mathcal{P}$ and $\mathcal{P}_\infty$, respectively, with homogeneous utility functions. To prove the theorem, we compare the utilities of the optimal utilitarian and envy-free schedules of these problems.

THEOREM 3. *Let $\mathcal{P}=(n,\mathcal{U},T,\tau,\beta)$, $n\in\mathbb{N}_{\geq 1}$, $\tau>0$, $\beta>0$, and for each $i\in[n]$, denote $\delta_i=\text{CADENCE}_{(\beta,\tau)}(u_i)$ and $\rho_i=e^{-\beta(\delta_i+\tau)}$. Set $i^*\in\arg\max_{i\in[n]}\frac{u_i(\delta_i)}{1-\rho_i}$, and let $\rho=\rho_{i^*}$. Assume that $T=m(\delta_{i^*}+\tau)-\tau$ for some integer $m>n$ such that $\rho^n\geq n\rho^m$, let $\lambda\in\mathbb{N}$ be the largest integer such that $\rho\geq\frac{\lambda-1}{\lambda}$, and let $\psi=\frac{\gamma(\mathcal{U})}{\gamma(\mathcal{U})+n-1}$. Then if $n>\lambda$,*

$$\text{PoF}(\mathcal{P})\leq\frac{\lambda\psi\,(1-\rho^m)}{\rho^{n-\lambda}-\lambda\rho^m},\tag{5}$$

otherwise (if $n\leq\lambda$), $\text{PoF}(\mathcal{P})\leq\frac{n\psi(1-\rho^m)}{1-n\rho^m}$.

Proof. Let $\mathcal{P}_T$, $\mathcal{P}_\infty$, $\mathcal{P}_{T,i}$ and $\mathcal{P}_{\infty,i}$ be as above. Then

$$\begin{aligned}
OPT\left(\mathcal{P}_T^{EF}\right)&\geq\frac{1}{n\psi}\max_{i\in[n]}OPT\left(\mathcal{P}_{T,i}^{EF}\right) &(21)\\
&\geq\frac{\rho^{n-\lambda}-\lambda\rho^m}{\psi\lambda}\max_{i\in[n]}OPT\left(\mathcal{P}_{\infty,i}^{UC}\right) &(22)\\
&=\frac{\rho^{n-\lambda}-\lambda\rho^m}{\psi\lambda}OPT\left(\mathcal{P}_\infty^{UC}\right) &(23)\\
&=\frac{\rho^{n-\lambda}-\lambda\rho^m}{\psi\lambda(1-\rho^m)}OPT\left(\mathcal{P}_T^{UC}\right). &(24)
\end{aligned}$$

Ineq. (21) is due to Lemma 15, which we state and prove below, in Subsection I.1. In this lemma, we show that if there exists a schedule π that is envy-free for homogeneous agents with utility function u_i (for some $i\in[n]$), then we can convert it to an envy-free allocation for heterogeneous agents with utility functions $u_1,\ldots,u_n$ using the following procedure. Let all agents except agent i (who has utility function u_i) choose their favorite allocation from π in serial fashion (i.e., each one chooses their favorite allocation from the remaining ones). Then, for all $j\in[n]\setminus\{i\}$, scale the durations of the intervals in agent j's allocation by some factor a_j while keeping the start times fixed. We show that there exists such a scaling $a_1,\ldots,a_n$ for which the resulting allocation is envy-free by using Brouwer's Fixed Point Theorem.

Ineq. (22) is due to Lemma 16, which we prove below. In this lemma, we show how to convert an optimal infinite-horizon UC schedule for homogeneous agents to a finite-horizon EF schedule, with only a small loss of utility. To do this, we note that all of the intervals of an optimal utilitarian schedule have an identical duration (by Lemma 11). As the schedule is tight, all agents have an identical utility function and $\beta > 0$, the utilities take the values of an infinite geometric series. We use the results on partitions of geometric series (Appendix B) to show that these intervals can be partitioned into sets such that the total utility for the intervals in each set is approximately the same. The assumption that $\rho^n \geq n\rho^m$ is made so that the bounds on $OPT\left(\mathcal{P}_{T,i}^{EF}\right)$ in Lemma 16 will be meaningful.

Eqs. (23) and (24) are due to results that we used to prove Theorem 1. Eq. (23) is due to Corollary 3 (a corollary to Theorem 1) which states that the optimal schedule for $\mathcal{P}_{\infty}^{UC}$ is exactly the optimal schedule for $\mathcal{P}_{\infty,i}^{UC}$ for the $i \in [n]$ which produces the highest utility. That is, it is the schedule with the highest utility out of the optimal schedules for $\mathcal{P}_{\infty,1}^{UC}, \mathcal{P}_{\infty,2}^{UC}, \ldots, \mathcal{P}_{\infty,n}^{UC}$. Finally, Eq. (24) is due to Lemma 12, which precisely characterizes the loss of utility due to truncating the infinite horizon optimal utilitarian schedule at T.

The proof for $n \leq \lambda$ is similar and is omitted. □

I.1. Lemma 15 (Converting an EF schedule for homogeneous agents to an EF schedule for heterogeneous agents)

LEMMA 15. *Let* $\mathcal{P}_T = (n, \mathcal{U}, T, \tau, \beta)$*,* $n \in \mathbb{N}_{\geq 1}$*,* $0 \leq T \leq \infty$*,* $\tau > 0$*,* $\beta \geq 0$*, and* $\psi = \frac{\gamma(\mathcal{U})}{\gamma(\mathcal{U})+n-1}$*. For* $i \in [n]$ *denote* $\mathcal{P}_{T,i} = (n, u_i, T, \tau, \beta)$*. Then*

$$OPT\left(\mathcal{P}_T^{EF}\right) \geq \frac{1}{n\psi} \max_{i \in [n]} OPT\left(\mathcal{P}_{T,i}^{EF}\right).$$

Proof. We will show that $OPT\left(\mathcal{P}_T^{EF}\right) \geq \frac{1}{n\psi} OPT\left(\mathcal{P}_{T,i}^{EF}\right)$ holds for every $i \in [n]$. Without loss of generality, we show that it holds for $i = n$.

Let π be an optimal schedule for $\mathcal{P}_{T,n}^{EF}$, where $\pi_i = \{(s_{i,1}, d_{i,1}), \ldots, (s_{i,K_j}, d_{i,K_i})\}$ (for some $K_i \geq 1$) for all $i \in [n]$. Because π is envy-free, agent n is indifferent between $\pi_1, \ldots, \pi_n$ (i.e., $u_n(\pi_i) = u_n(\pi_j)$ for all $i, j \in [n]$). Let $\pi_i^a = \{(s_{i,1}, a\, d_{i,1}), \ldots, (s_{i,K_i}, a\, d_{i,K_i})\}$ correspond to allocation π_i where all the interval durations have been scaled by $a \geq 0$. Now reassign $\pi_1, \ldots, \pi_n$ to the agents as follows: Set $\chi(1) = \arg\max_{i \in [n]}\{u_1(\pi_i)\}$, $\chi(2) = \arg\max_{i \in [n] \setminus \{\chi(1)\}}\{u_2(\pi_i)\}$, and so on. In other words, agent 1 chooses their favorite allocation, agent 2 chooses their favorite from the remaining allocations, and so on. Assume w.l.o.g. that $\chi(i) = i$ for all $i \in [n]$. Each agent (weakly) prefers their allocation to π_n. Let

$$y_i = \min_{y \geq 0} \{y : u_i(\pi_i^y) \geq u_i(\pi_n)\}$$

be the smallest possible scaling y of π_i so that agent i does not feel envy w.r.t. π_n. As agent i weakly prefers π_i to π_n, $y_i \le 1$.

For each $i \in \{1, \ldots, n-1\}$, define $f_i : [0,1] \to [0,1]^{n-1}$ coordinate-wise as follows:

$$f_i(a_i)[j] = \max_{y \in [0,1]} \left\{ y : u_i(\pi_i^{a_i}) \ge u_i(\pi_j^{y}) \right\}.$$

In simple terms, if π_i is scaled down by a_i, $f_i(a_i)[j]$ is the maximal amount that π_j can be scaled down such that agent i will feel no envy with respect to agent j's allocation.

Let $\vec{a} = (a_1, \ldots, a_{n-1})$ be the vector of scaling factors, where a_i is the i^{th} coordinate. Define the continuous function $f : [0,1]^{n-1} \to [0,1]^{n-1}$ coordinate-wise as follows:

$$f(\vec{a})[j] = \begin{cases} a_j & \text{if } y_j \le a_j \le \min_{i \in [n]} f_i(a_i)[j], \\ \min_{i \in [n-1]} f_i(a_i)[j] & \text{if } a_j > \min_{i \in [n]} f_i(a_i)[j], \\ y_j & \text{otherwise.} \end{cases}$$

Note that $f(\vec{a}) = \vec{a}$ if and only if the allocation that results from scaling π by $\vec{a}$ is envy-free. By Brouwer's Fixed Point Theorem, as f is a continuous function from a convex compact subset ($[0,1]^{n-1}$) to itself, it has a fixed point. Therefore there exists an envy-free allocation in which agent n is allocated π_n. For a fixed point $\vec{a}$,

$$\begin{aligned} OPT\left(\mathcal{P}_T^{EF}\right) &= \sum_{i \in [n]} u_i(\pi_i^{a_i}) \\ &\ge \sum_{i \in [n]} u_i(\pi_n) && (25) \\ &\ge u_n(\pi_n) + (n-1)\frac{u_n(\pi_n)}{\gamma(\mathcal{U})} && (26) \\ &= \left(\frac{\gamma(\mathcal{U}) + n - 1}{n\,\gamma(\mathcal{U})}\right) OPT\left(\mathcal{P}_{T,n}^{EF}\right), && (27) \end{aligned}$$

where Ineq. (25) is because the allocation is envy-free, Ineq. (26) is because for each agent $i \ne n$, $u_i(\pi_n) \ge \frac{u_n(\pi_n)}{\gamma(\mathcal{U})}$ (from the definition of heterogeneity), and Eq. (27) is because π was chosen to be an optimal schedule for $\mathcal{P}_{T,n}^{EF}$. □

I.2. Lemma 16 (Converting an optimal infinite-horizon UC schedule for homogeneous agents to a finite-horizon EF schedule)

LEMMA 16. *Let $\mathcal{P}_{\infty,u} = (n, u, \infty, \tau, \beta)$, $n \in \mathbb{N}_{\ge 1}$, $\tau > 0$, $\beta > 0$, $\delta = \text{CADENCE}_{(\beta,\tau)}(u)$, and $\rho = e^{-\beta(\delta+\tau)}$. Let $\lambda \in \mathbb{N}$ be the largest integer such that $\rho \ge \frac{\lambda-1}{\lambda}$. Let $m > n$ be some integer, set $T = m(\delta+\tau) - \tau$, and let $\mathcal{P}_u = (n, u, T, \tau, \beta)$. Then:*

(i) If $n \le \lambda$, $OPT\left(\mathcal{P}_u^{EF}\right) \ge (1 - n\rho^m)\, OPT\left(\mathcal{P}_{\infty,u}^{UC}\right)$.

(ii) If $n > \lambda$, $OPT\left(\mathcal{P}_u^{EF}\right) \ge \frac{n\left(\rho^{n-\lambda} - \lambda\rho^m\right)}{\lambda} OPT\left(\mathcal{P}_{\infty,u}^{UC}\right)$.

Proof. Let π be an optimal schedule for $\mathcal{P}^{UC}_{\infty,u}$, where all I_k have the same duration δ (by Lemma 11, such a schedule exists). Then $u(I_k) = \rho^{k-1}u(\delta)$ for all $k \geq 1$ which constitute a geometric series: $a, ar, ar^2, \ldots$, where $a = u(\delta)$ and $r = \rho$. We chose λ to be the largest integer such that $\rho \geq \frac{\lambda-1}{\lambda}$. By Lemma 11, $OPT\left(\mathcal{P}^{UC}_{\infty,u}\right) = \frac{u(\delta)}{1-\rho}$.

By Lemma 2, if $n \leq \lambda$, we can allocate the intervals to the agents such that each agent has utility at least $\frac{u(\delta)\left(1-n\rho^{m+1}\right)}{n(1-\rho)}$, and so $OPT\left(\mathcal{P}^{EF}_u\right) \geq \frac{u(\delta)\left(1-n\rho^{m+1}\right)}{1-\rho}$. If $n > \lambda$, we can allocate the intervals to the agents such that each agent has utility at least $\frac{u(\delta)\left(\rho^{n-\lambda}-n\rho^{m+1}\right)}{\lambda(1-\rho)}$, and so $OPT\left(\mathcal{P}^{EF}_u\right) \geq \frac{nu(\delta)\left(\rho^{n-\lambda}-n\rho^{m+1}\right)}{\lambda(1-\rho)}$, also by Lemma 2. □

Appendix J: Proof of Corollary 1

COROLLARY 1. *Let $\mathcal{P} = (n, \mathcal{U}, \infty, \tau, \beta)$, $n \in \mathbb{N}_{\geq 1}$, $\tau > 0$, $\beta > 0$, and define ρ, λ, and ψ as in Theorem 3. Then, if $n > \lambda$,*

$$\mathrm{PoF}(\mathcal{P}) \leq \frac{\lambda\psi}{\rho^{n-\lambda}},$$

otherwise $\mathrm{PoF}(\mathcal{P}) \leq n\psi$.

The following simple observation shows that it is possible to decrease the total utility of a given schedule by an arbitrarily small amount for a single agent; this is due to the continuity of the utility functions.

OBSERVATION 2. *Let $\mathcal{P} = (n, \mathcal{U}, T, \tau, \beta)$, let π be a feasible schedule for $\mathcal{P}$, and let $\mathcal{I}(\pi) = \{I_k\}_{k\geq 1}$ where $I_k = (s_k, d_k)$ for all $k \geq 1$. Then, for any $i \in [n]$ and any $0 \leq \epsilon \leq u_i(\pi, \mathcal{P})$ there exists a schedule π' with $\mathcal{I}(\pi') = \{I'_k\}_{k\geq 1}$ where $I'_k = (s_k, d'_k)$, $d'_k \leq d_k$ for all $k \geq 1$, such that $u_i(\pi', \mathcal{P}) = u_i(\pi, \mathcal{P}) - \epsilon$ and $u_j(\pi', \mathcal{P}) = u_j(\pi, \mathcal{P})$ for $j \neq i$.*

Proof of Corollary 1. Let $\mathcal{P}_\infty = (n, \mathcal{U}, \infty, \tau, \beta)$ and for any $T \geq 0$, let $\mathcal{P}_T = (n, \mathcal{U}, T, \tau, \beta)$. From Theorem 1, $\lim_{T\to\infty} OPT(\mathcal{P}^{UC}_T) = OPT(\mathcal{P}^{UC}_\infty)$.

Let π be an optimal schedule for $\mathcal{P}^{EF}_\infty$. Choose any $\epsilon > 0$, and set $\epsilon' = \frac{\epsilon}{n}$. Similarly to Theorem 2 (second bullet point), let $T = T_{\epsilon'}$ be such that for $T > T_{\epsilon'}$, for any agent i,

$$u_i(\pi^T) \geq u_i(\pi) - \epsilon/n,$$

where π^T is the truncation of π to include only those intervals that end at or before time T. Leveraging Observation 2, we construct an envy-free schedule for $\mathcal{P}^{EF}_T$ from π^T. As any agent's utility in π^T is at most ϵ/n less than its utility in π, we can construct a modified schedule $\tilde{\pi}^T$ such that the following hold: $s_k + d_k \leq T$ for all $I_k \in \mathcal{I}(\tilde{\pi}^T)$, $u_i(\tilde{\pi}^T) \geq u_i(\pi^T) - \epsilon/n$ for all $i \in [n]$, and $u_i(\tilde{\pi}^T) = u_j(\tilde{\pi}^T)$ for all $i, j \in [n]$. It follows that $\sum u_i(\tilde{\pi}^T) \geq \sum u_i(\pi) - \epsilon$, and so $OPT(\mathcal{P}^{EF}_{T'}) \geq OPT(\mathcal{P}^{EF}_\infty) - \epsilon$ for all $T' \geq T$. As ϵ is arbitrary, $\lim_{T\to\infty} OPT(\mathcal{P}^{EF}_T) = OPT(\mathcal{P}^{EF}_\infty)$.

Finally, since $PoF(\mathcal{P}_T) = \dfrac{OPT(\mathcal{P}^{UC}_T)}{OPT(\mathcal{P}^{EF}_T)}$, and $OPT(\mathcal{P}^{EF}_\infty) > 0$, we can take the limit of the ratio to get $\lim_{T\to\infty} PoF(\mathcal{P}_T) = PoF(\mathcal{P}_\infty)$, by taking the limit as $m \to \infty$ on both the LHS and RHS of Eq. (5) to complete the proof. □

Appendix K: Proof of Theorem 4

To prove Theorem 4, we prove the upper bound (Ineq. (6)) and give a matching lower bound. To prove the upper bound, we explicitly construct an envy-free schedule, similarly to the proof of Theorem 3. In this case, however, we allocate the resource more efficiently using the push right operation of Algorithm 1, thereby reducing the time the resource is left idle (for the most part; Figure 2a shows that the bound of Corollary 1 is sometimes tighter, showing that the schedule constructed in the proof of Theorem 3 is sometimes more efficient). The upper bound is shown by explicitly computing the utility of the envy-free schedule for homogeneous agents. The schedule is not optimal because: (i) we shorten some of the intervals to facilitate the exact computation of the utility; and (ii) the utility functions are heterogeneous. The lower bound is established by giving a family of instances for which there is no loss of utility from either source. The proofs for the upper and lower bounds are given in Subsections K.1 and K.2 respectively.

Theorem 4. *Let $\mathcal{P} = (n, \mathcal{U}, \infty, \tau, \beta)$, $n \in \mathbb{N}_{\geq 1}$, $\tau > 0$, $\beta > 0$, and for each $i \in [n]$, denote $\delta_i = \textsc{Cadence}_{(\beta,\tau)}(u_i)$ and $\rho_i = e^{-\beta(\delta_i+\tau)}$. Set $i^* \in \arg\max_{i\in[n]} \frac{u_i(\delta_i)}{1-\rho_i}$, and let $\rho = \rho_{i^*}$. Let $\lambda \in \mathbb{N}$ be the largest integer such that $\rho \geq \frac{\lambda-1}{\lambda}$, and let $\psi = \frac{\gamma(\mathcal{U})}{\gamma(\mathcal{U})+n-1}$. If $\lambda \geq n$ then $\mathrm{PoF}(\mathcal{P}) \leq n\psi$, otherwise*

$$\mathrm{PoF}(\mathcal{P}) \leq \frac{\psi u_{i^*}(\delta_{i^*})}{(1-\rho)\, u_{i^*}\left(\Phi^{n-\lambda}_{u_{i^*},\beta,\tau}(\delta_{i^*})\right)}. \tag{6}$$

Furthermore, there exists a family of instances for which this bound is tight.

Proof. We first prove the upper bound. For each $i \in [n]$, define $\mathcal{P}_{\infty,i} = (n, u_i, \infty, \tau, \beta)$. Without loss of generality, assume that $1 \in \arg\max_{i\in[n]} \frac{u_i(\delta_i)}{1-e^{-\beta(\delta_i+\tau)}}$ (this is because we can reorder the utility functions in $\mathcal{U}$). Then

$$OPT(\mathcal{P}^{EF}_{\infty}) \geq \frac{1}{n}\left(1 + \frac{n-1}{\gamma(\mathcal{U})}\right) OPT(\mathcal{P}^{EF}_{\infty,1}) \tag{28}$$

$$\geq \frac{(1-\rho)\, u_1\left(\Phi^{n-\lambda}_{u_1,\beta,\tau}(\delta)\right) OPT(\mathcal{P}^{UC}_{\infty,1})}{u_1(\delta)}\left(1 + \frac{n-1}{\gamma(\mathcal{U})}\right) \tag{29}$$

$$= \frac{(1-\rho)\, u_1\left(\Phi^{n-\lambda}_{u_1,\beta,\tau}(\delta)\right) OPT(\mathcal{P}^{UC}_{\infty})}{u_1(\delta)}\left(1 + \frac{n-1}{\gamma(\mathcal{U})}\right), \tag{30}$$

where Ineq. (28) is due to Lemma 15, which was used in the proof of Theorem 3, Ineq. (29) is due to Lemma 17, which we prove below, and Eq. (30) is due to Corollary 3 (since we assume that $1 \in \arg\max_{i\in[n]} OPT\left(\mathcal{P}^{UC}_{\infty,i}\right)$). Ineq. (6) is then obtained by rearranging Eq. (30).

We prove the lower bound in Lemma 20 below, which describes a family of instances for which these bounds are tight. □

K.1. Proof of the Upper Bound of Theorem 4

LEMMA 17. *Let $\mathcal{P}_{\infty,u} = (n, u, \infty, \tau, \beta)$, $\beta > 0$, and let $\delta = \text{CADENCE}_{(\beta,\tau)}(u)$, and $\rho = e^{-\beta(\delta+\tau)}$. Let λ be the largest integer in $[n]$ such that $\rho \geq \frac{\lambda-1}{\lambda}$. If $\lambda = n$ then* $\text{PoF}(\mathcal{P}_{\infty,u}) = 1$, *otherwise*

$$OPT(\mathcal{P}^{EF}_{\infty,u}) \geq \frac{n\,(1-\rho)\,u\left(\Phi^{n-\lambda}_{u,\beta,\tau}(\delta)\right) OPT(\mathcal{P}^{UC}_{\infty,u})}{u(\delta)}.$$

Proof. If $\lambda = n$, the result follows from Corollary 1 (by substituting $\gamma(\mathcal{U}) = 1$). For the rest of the proof, assume that $\lambda \leq n-1$. We construct an envy-free schedule similarly to Algorithm 1 for the case $T = \infty$ (i.e., there is no truncation), except that we (possibly) reduce some of the interval durations so as to accurately compute the utilities. By Lemma 11 there is an optimal utilitarian schedule for $\mathcal{P}_u$ which consists of intervals $\{I_k\}_{k\geq 1}$ where $I_k = (s_k, d_k)$, $s_j = (j-1)(\delta+\tau)$, and $d_k = \delta$ for all $j \geq 1$. Furthermore, $\{u(I_k)\}_{k\geq 1}$ is a geometric series where $u(I_k) = \rho^{k-1}u(\delta)$ for all $k \geq 1$. Similarly to Algorithm 1, we allocate these intervals greedily (using Algorithm 2) to agents $1, \ldots, \lambda$, where λ is defined in the lemma statement. By Corollary 2, there exist allocations π_i such that $u_i(\pi_i) = \frac{u(\delta)}{\lambda(1-\rho)}$ for agents $i \in \{1, \ldots, \lambda\}$.

As $1 - \rho \leq \frac{1}{\lambda}$, it holds that $\frac{u(\delta)}{\lambda(1-\rho)} \geq u(\delta)$. For each agent $i \in \{1, \ldots, \lambda\}$, construct a new allocation π'_i such that $u_i(\pi'_i) = u(\delta)$ should it just be $u(\pi'_i)$ on the LHS? It seems we switch to just writing u here and u with subscripts? , by keeping the start times fixed and only (possibly) reducing the durations. There exists such an allocation π'_i by Observation 2. Denote the intervals of π' by $\{I'_j = (s_j, d'_j)\}_{j\geq 1}$, where we keep $s_j = (j-1)(\delta+\tau)$ for all $j \geq 1$. We now extend $\pi' = \{\pi'_1, \ldots, \pi'_\lambda\}$ to an envy-free schedule π^{EF} as in the first loop of Algorithm 1. First, for $i \in [1, \lambda]$ we set $\pi^{EF}_i = \pi'_i$, and initialize $\pi^{EF} = \{\pi^{EF}_1, \ldots, \pi^{EF}_\lambda\}$. Then, for $i \in [\lambda+1, n]$: (i) let $d_i = y$ such that $u(y) = e^{-\beta(y+\tau)}u_{i-1}(\pi^{EF})$; (ii) let π^{EF}_i include (only) the interval $(0, d_i)$; and (iii) let $\pi^{EF} = \{\pi^{EF}_i\} \cup \pi^{EF} \triangleright (d_i + \tau)$.

Recall the definition of Φ given by $\Phi_{u,\beta,\tau}(t) := \{y \geq 0 : u(y) = e^{-\beta(y+\tau)}u(t)\}$. It is straightforward to verify that in each iteration of the loop ($i = \lambda+1, \ldots, n$) the length of the first interval, d_i, is $\Phi^{i-\lambda}_{u,\beta,\tau}(\delta)$. To see this, note that for $i = \lambda+1$ we have $u_{i-1}(\pi^{EF}) = u_\lambda(\pi') = u(\delta)$ and hence

$$d_{\lambda+1} = y : u(y) = e^{-\beta(y+\tau)}u(\delta) = \Phi_{u,\beta,\tau}(\delta),$$

and we continue iteratively.

Therefore, at the end of the procedure, the length of the first interval is $\Phi^{n-\lambda}_{u,\beta,\tau}(u(\delta))$, and

$$OPT(\mathcal{P}^{EF}_{\infty,u}) \geq \sum_{i\in[n]} u_i(\pi^{EF}_i) = nu\left(\Phi^{n-\lambda}_{u,\beta,\tau}(\delta)\right). \tag{31}$$

From Lemma 11, $OPT(\mathcal{P}^{UC}_{\infty,u}) = \frac{u(\delta)}{1-\rho}$ and therefore

$$\frac{OPT(\mathcal{P}^{EF}_{\infty,u})}{OPT(\mathcal{P}^{UC}_{\infty,u})} \geq \frac{nu\left(\Phi^{n-\lambda}_{u,\beta,\tau}(\delta)\right)(1-\rho)}{u(\delta)}, \tag{32}$$

completing the proof. □

K.2. Proof of the Lower Bound of Theorem 4

To establish the lower bound, we will use the following two problem instances, Instances 2 and 3. Instance 3 is a special case of Instance 2 with homogeneous utilities.

INSTANCE 2. *For any $\tau > 0$, $\beta > 0$, $\gamma \geq 1$, and integer $\lambda \geq 1$ such that $\frac{\log\lambda - \log(\lambda-1)}{\beta} > \tau$, set $c_1 = \frac{\log(\lambda+1)-\log\lambda}{\log\lambda-\log(\lambda-1)-\beta\tau}$, and set $c_i = \frac{c_1}{\gamma}$ for $i \geq 2$. For $i \geq 1$, let $u_i : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be piecewise linear concave such that $u_i(t) = c_i t$ for $0 \leq t \leq \eta$ and $u(t) = c_i \eta$ for $t > \eta$, where $\eta = \frac{\log\lambda-\log(\lambda-1)}{\beta} - \tau$. For any $n \geq \lambda$, let $\mathcal{P}_n = (n, \mathcal{U}_n, \infty, \tau, \beta)$ where $\mathcal{U}_n = \{u_1, u_2, \ldots, u_n\}$.*

INSTANCE 3. *For any $\tau > 0$, $\beta > 0$, and integer $\lambda \geq 1$ such that $\frac{\log\lambda-\log(\lambda-1)}{\beta} > \tau$, set $c = \frac{\log(\lambda+1)-\log\lambda}{\log\lambda-\log(\lambda-1)-\beta\tau}$, and let $u : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be piecewise linear concave such that $u(t) = ct$ for $0 \leq t \leq \eta$ and $u(t) = c\eta$ for $t > \eta$, where $\eta = \frac{\log\lambda-\log(\lambda-1)}{\beta} - \tau$. For $n \geq \lambda$, let $\dot{\mathcal{P}}_n = (n, u, \infty, \tau, \beta)$.*

Note that Instance 3 can be obtained from Instance 2 by setting $\gamma = 1$. We will show that Instance 3 is tight for Theorem 4. First, in Lemma 18, we compute the utility of the optimal envy-free schedule for Instance 3. Lemma 19 adapts this result to Instance 2. Finally, Lemma 20 uses these results to show that the bound of Theorem 4 is tight.

LEMMA 18. *For any $\tau > 0$, $\beta > 0$, and integer $\lambda \geq 1$ such that $\frac{\log\lambda-\log(\lambda-1)}{\beta} > \tau$, let u be as in Instance 3, let $\delta = \text{CADENCE}_{(\beta,\tau)}(u)$, and let $\rho = e^{-\beta(\delta+\tau)}$. For any $n \geq \lambda$, let $\mathcal{P}_n = (n, u, \infty, \tau, \beta)$. Then*

$$OPT\left(\mathcal{P}_n^{EF}\right) = nu\left(\Phi_{u,\beta,\tau}^{n-\lambda}(\delta)\right).$$

Proof. It is easy to verify that $\rho = \frac{\lambda-1}{\lambda}$. Consider the construction of the envy-free schedule in the proof of Lemma 17, and note that when $\rho = \frac{\lambda-1}{\lambda}$, Ineq. (31) holds with equality. Therefore Ineq. (32) also holds with equality using $OPT(\mathcal{P}_{\infty,u}^{UC}) = \frac{u(\delta)}{1-\rho}$. □

LEMMA 19. *For any $\tau > 0$, $\beta > 0$, $\gamma \geq 1$, and integer $\lambda \geq 1$ such that $\frac{\log\lambda-\log(\lambda-1)}{\beta} > \tau$, let $\mathcal{P}_n$ and $\dot{\mathcal{P}}_n$ be as in Instances 2 and 3, respectively. For any $n \geq 1$, let $\psi_n = \frac{\gamma}{\gamma+n-1}$, then*

$$OPT\left(\mathcal{P}_n^{EF}\right) = n\psi_n OPT\left(\dot{\mathcal{P}}_n^{EF}\right).$$

Proof. As the utility functions of Instance 2 are all scalar multiples of the utility function of Instance 3, a schedule π is envy-free for $\dot{\mathcal{P}}_n$ if and only if it is envy-free for $\mathcal{P}_n$, and furthermore a schedule is optimal for $\dot{\mathcal{P}}_n^{EF}$ if and only if it is optimal for $\mathcal{P}_n^{EF}$. Let π be an optimal envy-free schedule for $\mathcal{P}_n$. Assume that each agent receives utility z from π in $\mathcal{P}_n$, for a total utility of nz. Considering π as a schedule for $\dot{\mathcal{P}}_n$, agent 1 receives utility z, while the other $(n-1)$ agents receive utility $\frac{z}{\gamma}$, for a total utility of $z\left(1 + (n-1)\frac{1}{\gamma}\right)$. □

LEMMA 20. *For any $\tau > 0$, $\beta > 0$, $\gamma \geq 1$, and integer $\lambda \geq 1$ such that $\frac{\log\lambda-\log(\lambda-1)}{\beta} > \tau$, there exists an infinite sequence of utility functions $\mathcal{U}_\infty = \{u_1, u_2, \ldots\}$, where $\mathcal{U}_n := \{u_1, \ldots, u_n\}$ for $n \geq 1$, such that:*

(i) $1 \in \arg\max_{i\geq 1} \frac{u_i(\delta_i)}{1-e^{-\beta(\delta_i+\tau)}}$;

(ii) $e^{-\beta(\delta_1+\tau)} = \frac{\lambda-1}{\lambda}$;

(iii) *for any* $n \in \mathbb{N}_{\geq 2}$, $\gamma(\mathcal{U}_n) = \gamma$;

(iv) *setting* $\mathcal{P}_n = (n, \mathcal{U}_n, \infty, \tau, \beta)$,

- *if* $n \leq \lambda$, $\mathrm{PoF}(\mathcal{P}_n) = n\,\psi_n$,
- *if* $n > \lambda$,

$$\mathrm{PoF}(\mathcal{P}_n) = \frac{\psi_n u_1(\delta_1)}{(1-\rho)\, u_1\left(\Phi^{n-\lambda}_{u_1,\beta,\tau}(\delta_1)\right)},$$

where $\delta_i = \textsc{Cadence}_{(\beta,\tau)}(u_i)$, $\rho = \frac{\lambda-1}{\lambda}$, *and* $\psi_n = \frac{\gamma(\mathcal{U}_n)}{\gamma(\mathcal{U}_n)+n-1}$.

Proof. To prove the lemma, we show that Conditions (i)–(iv) hold for Instance 2. Conditions (i)–(iii) are straightforward to verify; it remains to show that Condition (iv) holds. To do so, we use Instance 3 (a special case of Instance 2, obtained by setting $\gamma = 1$). As noted in the proof of Lemma 19, a schedule is envy-free for Instance 2 if and only if it is envy-free for Instance 3. For any $n \geq \lambda$, let $\mathcal{P}_n$ and $\dot{\mathcal{P}}_n$ be as in Instances 2 and 3, respectively.

From Lemma 11, $OPT(\mathcal{P}^{UC}_n) = \frac{u_1(\delta_1)}{1-\rho}$. From Lemma 18, $OPT\,(\mathcal{P}^{EF}_n) = n u_1\left(\Phi^{n-\lambda}_{u_1,\beta,\tau}(\delta_1)\right)$. Therefore, by Lemma 19,

$$OPT\left(\dot{\mathcal{P}}^{EF}_n\right) = \frac{OPT\,(\mathcal{P}^{EF}_n)}{n\psi_n} = \frac{u_1\left(\Phi^{n-\lambda}_{u_1,\beta,\tau}(\delta_1)\right)}{\psi_n}.$$

Taking the ratio of $OPT(\mathcal{P}^{UC}_n)$ and $OPT\left(\dot{\mathcal{P}}^{EF}_n\right)$ completes the proof. □

Appendix L: Proof of Theorem 5

To prove Theorem 5, we first show, similarly to the proof of Theorem 1, that all of the intervals in an optimal utilitarian schedule are the same length. The differences are that here the result holds for all $T < \infty$, and the length of the interval is $\textsc{Bounded-Cadence}_{(T,\tau)}(u)$ rather than $\textsc{Cadence}_{(\beta,\tau)}(u)$. We then compute the utility of the optimal utilitarian schedule and of the envy-free schedule that is constructed from the optimal utilitarian schedule by deleting some intervals. We use these values to upper bound the price of fairness. For the lower bound, we give a family of instances such that the envy-free schedule constructed in this way is optimal.

Theorem 5. *Let* $\mathcal{P} = (n, u, T, \tau, 0)$, $n \in \mathbb{N}_{\geq 1}$, $0 < T < \infty$, *and* $\tau > 0$. *Let* $\delta = \textsc{Bounded-Cadence}_{(T,\tau)}(u)$, *and* $m = \frac{T+\tau}{\delta+\tau}$. *Then*

$$\mathrm{PoF}(\mathcal{P}) \leq \begin{cases} \frac{m}{m-(m \mod n)}, & \text{if } n \leq m, \\ \frac{T-(m-1)\tau}{T-(n-1)\tau}, & \text{if } n > m \text{ and } T > (n-1)\tau. \end{cases}$$

Otherwise, if $T \leq (n-1)\tau$, *then* $\mathrm{PoF}(\mathcal{P}) = \infty$.

Furthermore, for any $\tau > 0$ *and* $n \in \mathbb{N}_{\geq 1}$ *there exists a function* u *such that the above inequalities hold with equality for* $\mathcal{P}_T = (n, u, T, \tau, 0)$ *for countably many values of* T.

Proof. First note that $m = \frac{T+\tau}{\delta+\tau}$ is an integer since $(\delta+\tau)|(T+\tau)$ by definition of $\text{Bounded-Cadence}_{(T,\tau)}(u)$ (Definition 3). The proofs for the upper and lower bounds are given in Subsection L.2 and L.3, respectively. We give a short summary of the proof.

For the upper bound, note that if $T \le (n-1)\tau$, then no envy-free schedule exists and therefore $\text{PoF}(\mathcal{P}) = \infty$. Lemma 25 establishes the bound for the cases where $n \le m$ and when $n > m$ and $T > (n-1)\tau$.

For the lower bound, we handle the first case (where $n \le m$) in Lemma 26 and the second (where $n > m$ and $T > (n-1)\tau$) in Lemma 27. □

L.1. Preliminary Results for the Proof of Theorem 5

The following lemma is analogous to Lemma 10, but for the finite horizon undiscounted case.

Lemma 21. *Suppose u is non-decreasing, concave, and differentiable. Then, for any $T > 0$ and $1 \le m \le T/\tau$, the solution $d_1 = \cdots = d_m = \frac{T-(m-1)\tau}{m}$ is optimal for the optimization problem*

$$\max_{d_1,\ldots,d_m \in \mathbb{R}_{\ge 0}} \left\{ \sum_{k=1}^{m} u(d_k) \; : \; \sum_{k=1}^{m} d_k = T - (m-1)\tau \right\}.$$

Proof. Suppose the schedule π with $\mathcal{I}(\pi) = \{I_k\}_{k=1}^m$ where $I_k = (s_k, d_k)$ is optimal but that $d_1 = \cdots = d_m$ does not hold. Without loss of generality, suppose $d_1 < d_2$ (since we can re-order the intervals without changing the total utility in the undiscounted case). Choose $0 < \epsilon < d_2 - d_1$ and define the modified schedule π^ϵ with $\mathcal{I}(\pi^\epsilon) = \{I_k^\epsilon\}_{k=1}^m$ where $I_k^\epsilon = (s_k^\epsilon, d_k^\epsilon)$, $d_1^\epsilon = d_1 + \epsilon$, $d_2^\epsilon = d_2 - \epsilon$, and $d_k^\epsilon = d_k$ for all $k = 3, \ldots, m$. Because u' is non-increasing and $d_1 < d_2$, it holds that $\int_{d_1}^{d_1+\epsilon} u'(t)dt \ge \int_{d_2-\epsilon}^{d_2} u'(t)dt$. Therefore,

$$\sum_{i\in[n]} u(\pi^\epsilon) - \sum_{i\in[n]} u(\pi) = \int_{d_1}^{d_1+\epsilon} u'(t)dt - \int_{d_2-\epsilon}^{d_2} u'(t)dt \ge 0,$$

and so $\sum_{i\in[n]} u(\pi^\epsilon) \ge \sum_{i\in[n]} u(\pi)$. A similar argument holds for all remaining intervals. Hence there is an optimal solution where all intervals have the same duration. □

Lemma 22. *Let $f : \mathbb{R}_{\ge 0} \to \mathbb{R}_{\ge 0}$ be continuous, non-decreasing, concave, and satisfy $f(0) = 0$. Then, $\frac{f(y)}{f(t)} \le \frac{y}{t}$ for all $0 < t < y$.*

Proof. We rewrite

$$\frac{f(y)}{f(t)} = \frac{f(y) - f(t) + f(t)}{f(t)} = 1 + \frac{f(y) - f(t)}{f(t)}.$$

This expression is maximized for any function f that maximizes $(f(y) - f(t))/f(t)$. Since f' is non-increasing by concavity of f, the numerator $f(y) - f(t) = \int_x^y f'(\xi)d\xi$ is maximized when f' is constant, and hence when f is linear. □

Lemma 23. *Let* $\mathcal{P}=(n,u,T,\tau,0)$*,* $0<T<\infty$*, and*

$$m^{UC}=\arg\max_{\xi\in\mathbb{N}_{>0}}\xi\,u\left(\frac{T-(\xi-1)\tau}{\xi}\right). \tag{33}$$

Then there exists an optimal schedule π *for* $\mathcal{P}^{UC}$ *such that* $|\mathcal{I}(\pi)|=m^{UC}$ *and* $d_k=$ Bounded-Cadence$_{(T,\tau)}(u)$ *for all* $I_k=(s_k,d_k)\in\mathcal{I}(\pi)$*.*

Proof. By Lemma 21, there is some $\xi\in\mathbb{N}_{>0}$ such that there is an optimal schedule π^{UC} for $\mathcal{P}$ with $|\mathcal{I}(\pi^{UC})|=\xi$. In particular, the optimal value of ξ is determined by

$$\max_{\xi\in\mathbb{N}_{>0}}\xi\,u\left(\frac{T-(\xi-1)\tau}{\xi}\right).$$

Furthermore, $d_k=\frac{T-(\xi-1)\tau}{\xi}$ for all $I_k=(s_k,d_k)\in\mathcal{I}(\pi^{UC})$.

Next recall that Bounded-Cadence$_{(T,\tau)}(u)$ corresponds to the optimization problem:

$$\max_{t:(t+\tau)|(T+\tau)}\frac{u(t)}{t+\tau}. \tag{34}$$

Do the change of variables $t=\frac{T+\tau}{\xi}-\tau$, and note that: (i) ξ is an integer if and only if $(t+\tau)\mid(T+\tau)$; and (ii) the optimal value of Problem (34) must be obtained for $t>0$, because $u(0)=0$, and any $t>0$ gives nonzero average utility. We then see that Problem (34) is equivalent to the LHS of

$$\arg\max_{\xi\in\mathbb{N}_{>0}}\frac{\xi}{T+\tau}u\left(\frac{T-(\xi-1)\tau}{\xi}\right)=\arg\max_{\xi\in\mathbb{N}_{>0}}\xi\,u\left(\frac{T-(\xi-1)\tau}{\xi}\right). \tag{35}$$

Therefore, Bounded-Cadence$_{(T,\tau)}(u)=\frac{T-(m^{UC}-1)\tau}{m^{UC}}$ is the duration of each interval in the optimal schedule π^{UC}. □

Lemma 24. *Suppose* u *is concave, then* $f(t)=t\,u\left(\frac{T-t\tau}{t}\right)$ *is concave in* t *for* $t>0$*.*

Proof. Let $g(s,t)=s\,u(t/s)$ be the perspective of u, which is concave by concavity of u. For $s>0$, the function $g(T-s\tau,\,s)=s\,u\left(\frac{T-s\tau}{s}\right)$ is concave in s, since the composition of a concave function with an affine mapping is concave. □

L.2. Proof of the Upper Bound of Theorem 5

Lemma 25. *Let* $\mathcal{P}=(n,u,T,\tau,0)$ *and* $0<T<\infty$*. Let* π^{UC} *be an optimal schedule for* $\mathcal{P}^{UC}$ *such that* $d_k=\delta$ *for some* $\delta>0$ *for all* $I_k=(s_k,d_k)\in\mathcal{I}(\pi^{UC})$*. Let* $m^{UC}=|\mathcal{I}(\pi^{UC})|$*. Then*

- *If* $m^{UC}\geq n$*, then* $\mathrm{PoF}(\mathcal{P})\leq\frac{m^{UC}}{m^{UC}-(m^{UC}\mod n)}$*.*
- *If* $m^{UC}<n$ *and* $T>(n-1)\tau$*, then* $\mathrm{PoF}(\mathcal{P})\leq\frac{T-(m_{UC}-1)\tau}{T-(n-1)\tau}$*.*

Proof. For the first bullet point, note that such a schedule π^{UC} (i.e., one for which $d_k = \delta$ for some $\delta > 0$ for all $I_k = (s_k, d_k) \in \mathcal{I}(\pi)$) exists by Lemma 23. Let $K = m^{UC} - (m^{UC} \mod n)$, and let $\pi^K \subseteq \mathcal{I}(\pi^{UC})$ be some (arbitrary) subset of $\mathcal{I}(\pi^{UC})$ such that $|\pi^K| = K$ and $\pi_1^K, \ldots, \pi_n^K$ all have cardinality $\frac{K}{n}$ (note that $\frac{K}{n}$ is an integer). Clearly, π^K is feasible for $\mathcal{P}$; furthermore, π^K is envy-free because all π_i^K have the same number of intervals, all with the same duration, and there is no discounting. It holds that $\sum_{i\in[n]} u_i(\pi^K) = Ku(\delta)$, whereas $OPT(\mathcal{P}_u^{UC}) = \sum_{i\in[n]} u_i(\pi^{UC}) = m^{UC}u(\delta)$. Noting that $\sum_{i\in[n]} u_i(\pi^K)$ is a lower bound on $OPT(\mathcal{P}_u^{EF})$ completes the proof.

For the second bullet point, consider the schedule π^{EF} where each π_i^{EF} consists of a single interval $I_i = (s_i, d_i)$ with $d_i = \frac{T-(n-1)\tau}{n}$. This schedule is feasible as $T > (n-1)\tau$, and the total utility derived from it is $\sum_{i\in[n]} u_i(\pi^{EF}) = n\,u\left(\frac{T-(n-1)\tau}{n}\right)$. The total utility for the optimal schedule is $OPT(\mathcal{P}^{UC}) = m_{UC} u\left(\frac{T-(m_{UC}-1)\tau}{m_{UC}}\right)$ by Lemma 23. Note that the function $f(t) = t\,u\left(\frac{T-t\tau}{t}\right)$ is concave by Lemma 24. Therefore,

$$\begin{aligned}\mathrm{PoF}(\mathcal{P}) &\le \frac{m_{UC}\,u\left((T-(m_{UC}-1)\tau)/m_{UC}\right)}{n\,u\left((T-(n-1)\tau)/n\right)}\\ &\le \frac{T-(m_{UC}-1)\,\tau}{T-(n-1)\,\tau},\end{aligned} \tag{36}$$

where Ineq. (36) follows by Lemma 22. ☐

L.3. Proof of the Lower Bound of Theorem 5

We show that the bounds of Theorem 5 are tight separately. We will use a similar family of instances for each: Instances 4 and 5.

Instance 4. *For any integers $\tau > 0$ and $n > 0$, let $u(t) = t$ for $0 \le t \le \frac{\tau}{n}$ and $u(t) = \frac{\tau}{n}$ for $t > \frac{\tau}{n}$. For $T > 0$, define $\mathcal{P}_T = (n, u, T, \tau, 0)$.*

Instance 5. *For any integers $\tau > 0$ and $n > 0$, let $u(t) = t$ for $0 \le t \le \tau$ and $u(t) = \tau$ for $t > \tau$. For $T > 0$, define $\mathcal{P}_T = (n, u, T, \tau, 0)$.*

Lemma 26. *For any $n > 0$ and $\tau > 0$, let u be as in Instance 4. For any integer $m \ge n$, let $T_m = \tau(m-1) + \frac{m\tau}{n}$. Denote $\mathcal{P}_m = (n, u, T_m, \tau, 0)$. Then for all $m \ge n$, $\mathrm{PoF}(\mathcal{P}_m) = \frac{m}{m-(m \mod n)}$.*

Proof. Fix $m \ge n$. Let $\delta_m := \textsc{Bounded-Cadence}_{(T_m,\tau)}(u)$. It is easy to verify that $\delta_m = \frac{T_m+\tau}{m} - \tau$ and therefore $m = \frac{T_m+\tau}{\delta_m+\tau}$. By Lemma 23, there exists an optimal schedule π for $\mathcal{P}_m^{UC}$ such that $|\mathcal{I}(\pi)| = m^{UC}$, where (substituting the value for T_m into Problem (33)):

$$m^{UC} = \operatorname*{arg\,max}_{\xi\in\mathbb{N}_{>0}} \xi\, u\left(\frac{\tau(m-\xi) + \frac{m\tau}{n}}{\xi}\right).$$

Let $f(\xi) = \xi\, u\left(\frac{\tau(m-\xi) + \frac{m\tau}{n}}{\xi}\right)$. For u as in Instance 4, it is easy to verify that

$$f(\xi) = \begin{cases} \frac{\xi\tau}{n} & \text{if } 0 < \xi \le m,\\ \frac{m\tau}{n} - (\xi - m)\tau & \text{if } m < \xi \le m\left(1+\frac{1}{n}\right) - 1.\end{cases}$$

Note that $\xi \leq m\left(1+\frac{1}{n}\right)-1$ must hold since $T_m = \tau\left(m\left(1+\frac{1}{n}\right)-1\right)$. Hence $f(\xi)$ is uniquely maximized at $\xi = m$, and so $m^{UC} = m$.

If $m \mod n = 0$, then the optimal schedule for $\mathcal{P}_m^{UC}$, is also optimal for $\mathcal{P}_m^{EF}$ since we can directly assign intervals to create an envy-free schedule. To complete the proof, assume that $m \mod n \neq 0$.

Note that by Lemma 21, $f(\xi)$ gives the total utility of the optimal schedule for $\mathcal{P}_m^{UC}$ among all schedules that have exactly ξ intervals. Let $g(\xi)$ denote the total utility of the optimal schedule for $\mathcal{P}_m^{EF}$, among all schedules that have exactly ξ intervals. Clearly, $f(\xi) \geq g(\xi)$ for all $\xi > 0$. Furthermore, $f(\xi) = g(\xi)$ for any ξ such that $\xi \mod n = 0$. As f is increasing in ξ for $\xi \leq m$ and decreasing in ξ for $\xi > m$, we see that g is maximized for some $\xi \in [m-(m \mod n), m+n-(m \mod n)]$. Denote this maximizer as m^{EF}, and let π^{EF} be an optimal envy-free schedule for $\mathcal{P}_m^{EF}$ for which $|\mathcal{I}(\pi^{EF})| = m^{EF}$. We consider two cases: (i) $m-(m \mod n) \leq m^{EF} < m+n-(m \mod n)$ and (ii) $m^{EF} = m+n-(m \mod n)$.

(i) If $m-(m \mod n) \leq m^{EF} < m+n-(m \mod n)$, by the pigeonhole principle, there must be some π_i^{EF} such that $|\mathcal{I}(\pi_i^{EF})| \leq \lfloor m/n \rfloor$. As the utility of each interval is at most $\frac{\tau}{n}$, and the schedule is envy-free, we must have $\sum_{i\in[n]} u_i(\pi^{EF}) \leq n\frac{\tau}{n}\lfloor m/n \rfloor = m-(m \mod n)\frac{\tau}{n}$ (i.e., if some agent has at most $\lfloor m/n \rfloor$ intervals, then by envy-freeness no agent can have more than that).

(ii) If $m^{EF} = m+n-(m \mod n)$, then $m^{EF} \geq m+1$ since $m \mod n \neq 0$. Therefore, as $f(\xi)$ is decreasing in ξ for $\xi > m$,

$$f(m^{EF}) \leq f(m+1) = m-n < m-(m \mod n).$$

It follows that $OPT(\mathcal{P}_m^{EF}) \leq \frac{(m-(m \mod n))\tau}{n}$. As we can allocate each agent $\lfloor m/n \rfloor$ intervals each of length $\frac{\tau}{n}$, it holds that $OPT(\mathcal{P}_m^{EF}) \geq \frac{(m-(m \mod n))\tau}{n}$; hence $OPT(\mathcal{P}_m^{EF}) = \frac{(m-(m \mod n))\tau}{n}$. As $OPT(\mathcal{P}_m^{UC}) = \frac{m\tau}{n}$, $\mathrm{PoF}(\mathcal{P}_m) = \frac{m}{m-(m \mod n)}$. □

Lemma 27. *For any $n > 0$ and $\tau > 0$, let u be as in Instance 5. For any integer $1 \leq m < n$, let $T_m = 2\tau(m-1)$. Denote $\mathcal{P}_m = (n, u, T_m, \tau, 0)$ and let $\delta = \text{Bounded-Cadence}_{(T_m,\tau)}(u)$. Then for all $1 \leq m < n$, $\mathrm{PoF}(\mathcal{P}_m) = \frac{T_m-(m-1)\tau}{T_m-(n-1)\tau}$.*

Proof.

Similarly to the proof of Lemma 26, we can show that: (i) there exists an optimal schedule π for $\mathcal{P}_m^{UC}$ such that $|\mathcal{I}(\pi)| = m$ and $d_k = \tau$ for all $I_k = (s_k, d_k) \in \mathcal{I}(\pi)$; and (ii) the optimal envy-free schedule consists of n intervals, each of duration $\frac{(2m-n)\tau}{n}$. Therefore, $OPT(\mathcal{P}_m^{UC}) = m\tau$ and $OPT(\mathcal{P}_m^{EF}) = (2m-n)\tau$ (where $2m-n > 0$ because $T_m > (n-1)\tau$). We conclude that

$$\mathrm{PoF}(\mathcal{P}_m) = \frac{m\tau}{(2m-n)\tau} = \frac{T_m-(m-1)\tau}{T_m-(n-1)\tau},$$

as required. □

Appendix M: Proof of Theorem 6

In order to prove Theorem 6, we require the following lemma, which is analogous to Lemma 11 for the case $\beta = 0$.

LEMMA 28. *Let $\mathcal{P}_u = (n, u, \infty, \tau, 0)$, and set $\delta = \text{CADENCE}_\tau(u)$. Set $I_k = (s_k, d_k)$ with $s_k = (k-1)(\delta + \tau)$ and $d_k = \delta$ for all $k \geq 1$. Then any $\pi = (\pi_1, \ldots, \pi_n)$ with $\mathcal{I}(\pi) = \{I_k\}_{k \geq 0}$ is an optimal schedule for $\mathcal{P}_u^{UC}$.*

Proof. We momentarily relax the constraint that m must be integer and consider:

$$m^* = \arg\max_{m \in \mathbb{R}_{>0}} m\, u\left(\frac{T-(m-1)\tau}{m}\right). \tag{37}$$

The change of variables

$$d(m,T) = \frac{T-(m-1)\tau}{m} \quad \text{and} \quad m(d,T) = \frac{T+\tau}{d+\tau},$$

gives two equivalent forms of Problem (37):

$$\max_{m \geq 0} m\, u\left(\frac{T-(m-1)\tau}{m}\right) \equiv (T+\tau) \max_{d \geq 0} \frac{u(d)}{d+\tau}.$$

Let

$$\delta = (T+\tau) \arg\max_{d \geq 0} \frac{u(d)}{d+\tau},$$

and note that δ does not depend on T (since $\arg\max$ is invariant under scaling). By the change of variables, for any $T > 0$ the optimal solution to Problem (37) is

$$m^*(T) = m(\delta, T) = \frac{T+\tau}{\delta+\tau}. \tag{38}$$

Now, the optimal average utility in Problem (37) when we are restricted to integer m is upper bounded by $\frac{m^*(T)u(\delta)}{T}$ (this is the optimal average utility without the integrality constraint). Since $m^*(T)\,\delta + (m^*(T)-1)\tau = T$ by assumption, it is immediate that a schedule with $\lfloor m^*(T) \rfloor$ intervals each of duration δ is feasible (where the resource is just left idle for any excess time). Hence, $\frac{\lfloor m^*(T)\rfloor u(\delta)}{T}$ is a lower bound on the optimal average utility when m is constrained to be integer.

By Eq. (38), we see that

$$\lim_{T \to \infty} \left[\frac{m^*(T)u(\delta)}{T} - \frac{\lfloor m^*(T)\rfloor u(\delta)}{T}\right] = 0,$$

i.e., the upper and lower bounds match as $T \to \infty$. Taking the limit as $T \to \infty$, we see that the optimal average utility is then

$$\lim_{T \to \infty} \frac{m^*(T)u(\delta)}{T} = \lim_{T \to \infty} \frac{T+\tau}{T} \frac{u(\delta)}{\delta+\tau} = \frac{u(\delta)}{\delta+\tau}.$$

The schedule where all intervals have duration δ attains the optimal average utility, and is thus an optimal schedule for $\mathcal{P}_u^{UC}$. □

Theorem 6. *Let* $\mathcal{P} = (n, \mathcal{U}, \infty, \tau, 0)$*,* $n \in \mathbb{N}_{\geq 1}$*,* $\tau > 0$*, and* $\psi = \frac{\gamma(\mathcal{U})}{\gamma(\mathcal{U})+n-1}$*. Then* $\mathrm{PoF}(\mathcal{P}) \leq n\psi$*. Furthermore, for any* $n \in \mathbb{N}_{\geq 1}$ *and* $\tau > 0$*, there exists a set of* n *utility functions* $\mathcal{U}$ *for which this bound is tight.*

Proof. For each $i \in [n]$, define $\mathcal{P}_i = (n, u_i, \infty, \tau, 0)$ and let $\delta_i = \textsc{Cadence}(u_i)$. By Lemma 28, there exists an optimal schedule π^i for $\mathcal{P}_i^{UC}$ where all intervals have duration δ_i. We construct an envy-free schedule from π^i by assigning the intervals in a round robin fashion, and therefore

$$OPT\left(\mathcal{P}_i^{EF}\right) = OPT\left(\mathcal{P}_i^{UC}\right). \tag{39}$$

We can then bound $OPT\left(\mathcal{P}^{EF}\right)$ as follows:

$$\begin{aligned} OPT\left(\mathcal{P}^{EF}\right) &\geq \max_{i \in [n]} \frac{OPT\left(\mathcal{P}_i^{EF}\right)}{n\psi} &(40)\\ &= \max_{i \in [n]} \frac{OPT\left(\mathcal{P}_i^{UC}\right)}{n\psi} &(41)\\ &= \frac{OPT\left(\mathcal{P}^{UC}\right)}{n\psi}, &(42)\end{aligned}$$

where Ineq. (40) is due to Lemma 15, Eq. (41) follows from Eq. (39), and Eq. (42) is due to Corollary 3. □

Appendix N: Proof of Propositions

N.1. Proof of Proposition 1

Proposition 1. *Fix* $n \in \mathbb{N}_{\geq 2}$ *and a utility function* u*. For* $\beta > 0$ *and* $\tau > 0$*, let* $\mathcal{P}_{(\beta,\tau)} = (n, u, \infty, \tau, \beta)$*. Then* $\mathrm{PoF}(\mathcal{P}_{(\beta,\tau)})$ *grows exponentially as a function of* β *and* τ*, i.e.,* $\mathrm{PoF}(\mathcal{P}_{(\beta,\tau)}) = e^{\Theta(\beta+\tau)}$*.*

Proof. We prove the result for β; the proof for τ is similar and omitted. Fix $n \geq 2$, $\tau > 0$, and u. For $\beta > 0$, denote $\mathcal{P}_\beta = (n, u, \infty, \tau, \beta)$. We show that $\mathrm{PoF}(\mathcal{P}_\beta)$ grows exponentially as a function of β (i.e., that $\mathrm{PoF}(\mathcal{P}_\beta) = e^{\Theta(\beta)}$). To establish this claim, we show that $\mathrm{PoF}(\mathcal{P}_\beta) = e^{O(\beta)}$ and $\mathrm{PoF}(\mathcal{P}_\beta) = e^{\Omega(\beta)}$.

(i) $\mathrm{PoF}(\mathcal{P}_\beta) = e^{O(\beta)}$**.** Let $u_\beta(t) := \int_0^t e^{-\beta x} \upsilon(x) dx$ for all $t \geq 0$ for $\beta > 0$. Fix some β_0 and let $\delta_0 = \textsc{Cadence}_{(\beta_0,\tau)}(u)$ and $\rho_0 = e^{-\beta(\delta_0+\tau)}$. Set $\beta \geq \beta_0$, and let $\delta = \textsc{Cadence}_{(\beta,\tau)}(u_\beta)$ and $\rho = e^{-\beta(\delta+\tau)}$. From Lemma 7 part *(ii)*, $\delta = \textsc{Cadence}_{(\beta,\tau)}(u_\beta)$ is increasing in β for $\beta \geq 1/\tau$. Therefore $\delta \geq \delta_0$; hence $\rho \leq \rho_0$. From Corollary 1, (as $\gamma(\mathcal{U}) = 1$, hence $\psi = \frac{1}{n}$), we then have

$$\mathrm{PoF}(\mathcal{P}_\beta) \leq \frac{\lambda}{n\rho^{n-\lambda}} \leq \frac{1}{\rho^n} = e^{\beta(\delta+\tau)}.$$

(ii) $\mathrm{PoF}(\mathcal{P}_\beta) = e^{\Omega(\beta)}$**.** Consider the optimal envy-free schedule, and rearrange the agents by the start time of the earliest interval allocated to them. As there are at least $n-1$ switches before the start of the first interval of agent n, and the maximum utility attainable from any schedule that begins at time t is $e^{-\beta t} OPT(\mathcal{P}_\beta^{UC})$, agent n receives at most $e^{-\beta(n-1)\tau} OPT(\mathcal{P}_\beta^{UC})$. Because the schedule is envy-free, no agent receives more utility than this, hence

$$\mathrm{PoF}(\mathcal{P}_\beta) \geq \frac{OPT(\mathcal{P}_\beta^{UC})}{n e^{-\beta(n-1)\tau} OPT(\mathcal{P}_\beta^{UC})} = \frac{e^{\beta(n-1)\tau}}{n},$$

as required. □

N.2. Proof of Proposition 2

We require a preliminary result, Lemma 29, in which we establish a property of the discounted reciprocal function $\Phi_{f,\beta,\tau}$.

LEMMA 29. *For any continuous and non-decreasing function $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, $\beta > 0$, $\tau > 0$, and $p \geq 1$,*

$$f\left(\Phi_{f,\beta,\tau}^p(t)\right) = e^{-\beta\left(\sum_{k=1}^p \Phi_{f,\beta,\tau}^k(t) + p\tau\right)} f(t),\ \forall t \geq 0.$$

Proof. The proof is by induction. For the base case, $p = 1$, note that $f(y) = e^{-\beta(y+\tau)} f(t)$ for $y = \Phi_{f,\beta,\tau}(t)$ by Definition 4; therefore

$$f\left(\Phi_{f,\beta,\tau}(t)\right) = e^{-\beta\left(\Phi_{f,\beta,\tau}(t)+\tau\right)} f(t). \tag{43}$$

For the induction step, assume that $f\left(\Phi_{f,\beta,\tau}^{p-1}(t)\right) = e^{-\beta\left(\sum_{k=1}^{p-1} \Phi_{f,\beta,\tau}^k(t) + (p-1)\tau\right)} f(t)$ for all $x \geq 0$. Then

$$\begin{aligned}
f\left(\Phi_{f,\beta,\tau}^p(t)\right) &= f\left(\Phi_{f,\beta,\tau}\left(\Phi_{f,\beta,\tau}^{p-1}(t)\right)\right) && (44)\\
&= e^{-\beta\left(\Phi_{f,\beta,\tau}\left(\Phi_{f,\beta,\tau}^{p-1}(t)\right)+\tau\right)} f\left(\Phi_{f,\beta,\tau}^{p-1}(t)\right) && (45)\\
&= e^{-\beta\left(\Phi_{f,\beta,\tau}^p(t)+\tau\right)} f\left(\Phi_{f,\beta,\tau}^{p-1}(t)\right) \\
&= e^{-\beta\left(\Phi_{f,\beta,\tau}^p(t)+\tau\right)} e^{-\beta\left(\sum_{k=1}^{p-1} \Phi_{f,\beta,\tau}^k(t)+(p-1)\tau\right)} f(t) && (46)\\
&= e^{-\beta\left(\sum_{k=1}^p \Phi_{f,\beta,\tau}^k(t)+p\tau\right)} f(t),
\end{aligned}$$

where Eq. (44) is due to Eq. (4), Eq. (45) is due to the base case (Eq. (43)), and Eq. (46) is due to the induction hypothesis. □

PROPOSITION 2. *For any $\tau > 0$, $\beta > 0$, and $\lambda \in \mathbb{N}_{\geq 1}$, there exists a family of instances $\{\mathcal{P}_n\}_{n=\lambda}^\infty$ where $\mathcal{P}_n = (n, u, \infty, \tau, \beta)$ such that for $n = \lambda$, $\mathrm{PoF}(\mathcal{P}_n) = 1$, and for $n > \lambda$, the price of fairness grows exponentially in n, i.e., $\mathrm{PoF}(\mathcal{P}_n) = \Omega(2^n)$.*

Proof. To prove the proposition, we use Instance 3. Recall that this instance is a special case of Instance 2, which we used to prove Lemma 20. The case $n=\lambda$ is immediate from Lemma 20 (part *(iv)*), as $n\psi_n = 1$. For $n>\lambda$,

$$\begin{aligned}\mathrm{PoF}(\mathcal{P}_n) &= \frac{\psi_n u(\delta)}{(1-\rho)\, u\left(\Phi^{n-\lambda}_{u,\beta,\tau}(\delta)\right)} &(47)\\ &= \frac{\psi_n u(\delta)}{(1-\rho)\, e^{-\beta\left(\sum_{k=1}^{n-\lambda}\Phi^k_{u,\beta,\tau}(x)+(n-\lambda)\tau\right)} u(\delta)} &(48)\\ &= \psi_n \lambda e^{\beta\left(\sum_{k=1}^{n-\lambda}\Phi^k_{u,\beta,\tau}(x)+(n-\lambda)\tau\right)},\end{aligned}$$

where $\rho = \frac{\lambda-1}{\lambda}$, Eq. (47) is due to Lemma 20, and Eq. (48) is due to Lemma 29. It is immediate that this expression grows exponentially in n due to the term $n\tau$ in the exponent and the fact that $\sum_{k=1}^{n-\lambda}\Phi^k_{u,\beta,\tau}(x)$ is always non-negative. □

N.3. Proof of Proposition 3

Proposition 3. *Let $\mathcal{P}_u = (n,u,\infty,\tau,\beta)$, $n\in\mathbb{N}_{\geq 1}$, $\tau>0$, $\beta\geq 0$, for arbitrary u. Let $\gamma\in\mathbb{R}_{\geq 1}$, and let $\mathcal{P}=(n,\mathcal{U},\infty,\tau,\beta)$, where $\mathcal{U}=\{u_1,\dots,u_n\}$, $u_1=u$, and u_i satisfies $\frac{u(t)}{\gamma}\leq u_i(t)\leq u(t)$ for all $i\in[2,n]$ and $t\geq 0$. Then*

$$\mathrm{PoF}(\mathcal{P}) \leq \min\{\gamma,n\}\cdot \mathrm{PoF}(\mathcal{P}_u).$$

Proof. First, it is clear that $OPT\,(\mathcal{P}^{UC}_u) = OPT\,(\mathcal{P}^{UC})$ by choice of $\mathcal{U}$. Since $\gamma(\mathcal{U})\leq\gamma$, Lemma 15 gives the bound

$$OPT\left(\mathcal{P}^{EF}\right) \geq \left(\frac{\gamma(\mathcal{U})+n-1}{n\,\gamma(\mathcal{U})}\right) OPT\left(\mathcal{P}^{EF}_u\right).$$

Therefore,

$$\mathrm{PoF}\left(\mathcal{P}\right) \leq \frac{n\,\gamma(\mathcal{U})}{\gamma(\mathcal{U})+n-1}\mathrm{PoF}\left(\mathcal{P}_u\right) \leq \frac{n\,\gamma}{\gamma+n-1}\mathrm{PoF}\left(\mathcal{P}_u\right).$$

Finally, it is straightforward to verify that $\frac{n\,\gamma}{\gamma+n-1}\leq\min\{\gamma,n\}$ for $\gamma\geq 1$ and $n\geq 1$. □

N.4. Proof of Proposition 4

We require a preliminary result, Lemma 30, in which we compute the cadence for Instance 1.

Lemma 30. *Let β, τ, and n be as in Instance 1. Let u_x be defined as in Instance 1. Then* $\textsc{Cadence}_{(\beta,\tau)}(u_x) = 1$ *for every $x\in[0.5,1]$.*

Proof. Let $g(t) = \frac{u_x(t)}{1-e^{-\beta(t+\tau)}}$. Taking the derivative for $t\in[0,1]$ gives:

$$\frac{dg(t)}{dt} = \frac{xt^{x-1}(1-e^{-\beta(t+\tau)}) - t^x\beta e^{-\beta(t+\tau)}}{\left(1-e^{-\beta(t+\tau)}\right)^2}.$$

Simplifying the numerator and using $\beta=0.1$, we get

$$xt^{x-1}(1-e^{-\beta(t+\tau)}) - t^x\beta e^{-\beta(t+\tau)} = xt^{x-1} - e^{-0.1(t+1)}(xt^{x-1}+0.1t^x).$$

We want to determine the sign of $xt^{x-1} - e^{-0.1(t+1)}(xt^{x-1} + 0.1t^x)$, so we first factor

$$xt^{x-1} - e^{-0.1(t+1)}(xt^{x-1} + 0.1t^x) = t^{x-1}\left(x - e^{-0.1(t+1)}(x + 0.1t)\right).$$

The function t^{x-1} is nonnegative for all $t \in [0,1]$, so we only need to check the sign of

$$\begin{aligned} x - e^{-0.1(t+1)}(x + 0.1t) =& (1 - e^{-0.1(t+1)})x - (0.1)t\, e^{-0.1(t+1)} \\ \geq& (1 - e^{-0.1(t+1)})(0.5) - (0.1)t\, e^{-0.1(t+1)} \\ =& 0.5 - e^{-0.1(t+1)}\left(0.5 + (0.1)t\right), \end{aligned}$$

where the inequality follows because $1 - e^{-0.1(t+1)} \geq 0$ for all $t \in [0,1]$. Taking the derivative, we can see that the function $e^{-0.1(t+1)}\left(0.5 + (0.1)t\right)$ is increasing in t for all $t \in [0,1]$, so

$$0.5 - e^{-0.1(t+1)}\left(0.5 + (0.1)t\right) \geq 0.5 - 0.6e^{-0.2} > 0.$$

We conclude that $x - e^{-0.1(t+1)}(x + 0.1t) > 0$ for all $x \in [0.5,1]$ and $t \in [0,1]$. It follows that $g(t)$ is increasing for all $t \in [0,1]$. As $g(t)$ is decreasing for $t > 1$, $g(t)$ is maximized at $t = 1$. □

PROPOSITION 4. *Let β, τ, and n be as in Instance 1. Then for every $x, x' \in [0.5,1]$ such that $x < x'$, it holds that* $\mathrm{PoF}(\mathcal{P}_x) < \mathrm{PoF}(\mathcal{P}_{x'})$.

Proof. By Lemma 30, for any $x \in [0.5,1]$, $\textsc{Cadence}_{(\beta,\tau)}(u_x) = 1$ and so

$$OPT(P_x^{UC}) = \frac{u_x(1)}{1 - e^{-\beta(\tau+1)}} = 4u_x(1) = 4.$$

Similarly to the proof of Lemma 17 we can show that $OPT(P_x^{EF}) = nu_x\left(\Phi^{n-\lambda}_{u_x,\beta,\tau}(1)\right)$, and so

$$\begin{aligned} \mathrm{PoF}(\mathcal{P}_x) &= \frac{4}{nu_x\left(\Phi^{n-\lambda}_{u_x,\beta,\tau}(1)\right)} \\ &= \frac{4}{ne^{-\beta\left(\sum_{k=1}^{n-\lambda}\Phi^k_{u_x,\beta,\tau}(x)+p\tau\right)}u_x(1)}, \end{aligned} \tag{49}$$

where Eq. (49) is due to Lemma 29. Then,

$$\frac{\mathrm{PoF}(\mathcal{P}_x)}{\mathrm{PoF}(\mathcal{P}_{x'})} = \frac{e^{-\beta\left(\sum_{k=1}^{n-\lambda}\Phi^k_{u_{x'},\beta,\tau}(1)+p\tau\right)}}{e^{-\beta\left(\sum_{k=1}^{n-\lambda}\Phi^k_{u_x,\beta,\tau}(1)+p\tau\right)}}.$$

It is easy to verify that u_x and $u_{x'}$ satisfy Lemma 9 for $c = 1$. Therefore, for every $p \geq 1$, $\Phi^p_{u_x,\beta,\tau}(1) \leq \Phi^p_{u_{x'},\beta,\tau}(1)$, and so $\frac{\mathrm{PoF}(\mathcal{P}_x)}{\mathrm{PoF}(\mathcal{P}_{x'})} \leq 1$. □

N.5. Proof of Proposition 5

Proposition 5. *For any $\tau > 0$, $T < \infty$, and utility function u, let $\mathcal{P}_n = (n, u, T, \tau, 0)$ for $n \in \mathbb{N}_{\geq 2}$. There exists some $m < \frac{T}{\tau} + 1$ such that* $\text{PoF}(\mathcal{P}_n) \leq 2$ *if $n < m$,* $\text{PoF}(\mathcal{P}_n) = \infty$ *if $n \geq \frac{T}{\tau} + 1$, and* $\text{PoF}(\mathcal{P}_n)$ *is increasing and convex (in the discrete sense) in n if $m \leq n < \frac{T}{\tau} + 1$.*

Proof. Let m be the number of intervals in the optimal schedule for $\mathcal{P}_n^{UC}$ (note that the number of intervals in the optimal schedule for $\mathcal{P}_n^{UC}$ is independent of n by Lemma 23). From Theorem 5, if $n \leq m$, then

$$\text{PoF}(\mathcal{P}_n) \leq \frac{m}{m - (m \mod n)}.$$

If $0 < n \leq \frac{m}{2}$, then

$$\frac{m}{m - (m \mod n)} \leq \frac{m}{m - n} \leq \frac{m}{m - \frac{m}{2}} = 2.$$

If $\frac{m}{2} < n \leq m$, then

$$\frac{m}{m - (m \mod n)} = \frac{m}{n} \leq 2.$$

If $m \leq n < \frac{T}{\tau} + 1$, let $\delta = \text{Bounded-Cadence}_{(T,\tau)}(u)$. Then $OPT\,(\mathcal{P}_n^{UC}) = m\,u(\delta)$ and $OPT\,(\mathcal{P}_n^{EF}) = nu\left(\frac{T-(n-1)\tau}{n}\right)$. Set

$$f(n) = \frac{OPT\,(\mathcal{P}_n^{UC})}{OPT\,(\mathcal{P}_n^{EF})} = \frac{m\,u(\delta)}{n\,u\left(\frac{T-(n-1)\tau}{n}\right)},$$

which is the price of fairness for n agents.

Define functions g and h by $g(x) = x\,u\left(\frac{T-(x-1)\tau}{x}\right)$ and $h(x) = m\,u(\delta)/x$, respectively. Then, $f = h \circ g$ is strictly convex since it is the composition of the strictly convex and decreasing function h and the strictly concave function g (on the domain where g is positive). By strict convexity of f, we have

$$\frac{1}{2}f(n+1) + \frac{1}{2}f(n-1) > f\left(\frac{1}{2}(n+1) + \frac{1}{2}(n-1)\right) = f(n),$$

which shows that $\text{PoF}(\mathcal{P}_{n+1}) - \text{PoF}(\mathcal{P}_n) > \text{PoF}(\mathcal{P}_n) - \text{PoF}(\mathcal{P}_{n-1})$. For $n = m$, $\text{PoF}(\mathcal{P}_n) = 1$. As the price of fairness cannot be less than 1 and $\text{PoF}(\mathcal{P}_n)$ is convex on $[m, \frac{T}{\tau} + 1)$, it must be increasing on this interval. □

N.6. Proof of Proposition 6

Proposition 6. *For any utility function u, the set of instances $\{\mathcal{P}_{n,\tau,T} = (n, u, T, \tau, 0), n \in \mathbb{N}_{>0}, 0 \leq \tau \leq T \leq \infty\}$ is recurrently optimal in τ and T and asymptotically recurrently optimal in n.*

We prove each part separately. Proposition 7 (below) shows that $\mathcal{P}_{n,\tau,T}$ is recurrently optimal in τ; Proposition 8 shows that $\mathcal{P}_{n,\tau,T}$ is recurrently optimal in T; and Proposition 7 shows that $\mathcal{P}_{n,\tau,T}$ is asymptotically recurrently optimal in n.

N.6.1. Recurrent optimality in τ

PROPOSITION 7. *Fix $n \in \mathbb{N}_{\geq 1}$, u, and $T > 0$, and define $\mathcal{P}_\tau = (n, u, T, \tau, 0)$ for $0 < \tau < T$. For any τ' such that $(n+1)\textsc{Cadence}_{\tau'}(u) + n\tau' < T$ for which $\text{PoF}(\mathcal{P}_{\tau'}) > 1$, there exists $\tau'' > \tau'$ for which $\text{PoF}(\mathcal{P}_{\tau''}) = 1$.*

Proof. We first consider the optimal schedule for $\mathcal{P}_\tau^{UC}$ for general $0 < \tau < T$. As in the proof of Theorem 6, we momentarily relax the constraint that m must be integer to obtain Problem (37) for which there are two equivalent forms:

$$\max_{m \geq 0} m\, u\left(\frac{T-(m-1)\tau}{m}\right) \equiv (T+\tau) \max_{d \geq 0} \frac{u(d)}{d+\tau}.$$

Next, let $\delta(\tau) = \textsc{Cadence}_\tau(u)$ (here we emphasize the dependence on τ) and let

$$m^*(\tau) = \frac{T+\tau}{\delta(\tau)+\tau}$$

be the optimal solution of Problem (37) (as a function of τ). The corresponding optimal solution of the integer-constrained problem

$$\max_{m \in \mathbb{N}_{>0}} m\, u\left(\frac{T-(m-1)\tau}{m}\right)$$

is then either $\lfloor m^*(\tau) \rfloor$ or $\lceil m^*(\tau) \rceil$ (the two integer neighbors of $m^*(\tau)$) by concavity of $m \to m\, u\left(\frac{T-(m-1)\tau}{m}\right)$. By Lemma 7, we recall that $\delta(\tau) = \textsc{Cadence}_\tau(u)$ is strictly increasing in τ. It follows that $m^*(\tau) = \frac{T+\tau}{\delta(\tau)+\tau}$ is strictly decreasing in τ.

For τ' as given, $(n+1)\textsc{Cadence}_{\tau'}(u) + n\tau' < T$ which is equivalent to

$$m^*(\tau') = \frac{T+\tau'}{\delta(\tau')+\tau'} > n+1,$$

and $\text{PoF}(\mathcal{P}_{\tau'}) > 1$ so $m^*(\tau') \mod n \neq 0$. Because $m^*(\tau)$ is strictly decreasing in τ, we can increase from τ' to some τ'' where $m^*(\tau'') \mod n = 0$. In this case, we can allocate the intervals of the optimal schedule equally, so that $\text{PoF}(\mathcal{P}_{\tau''}) = 1$. □

The restriction on τ' in Proposition 7 is because if $m < n$, it is not possible to obtain a price of fairness of one in all but degenerate cases.

N.6.2. Recurrent optimality in T. Flexibility in the time horizon in either direction (i.e., even if we are only allowed to increase or decrease the time horizon relative to some sufficiently large benchmark) allows us to guarantee that there is no loss in value due to the fairness constraint. We note that the cadence appears in the proposition below through $\textsc{Cadence}_\tau$ (which does not depend on T), and not $\textsc{Bounded-Cadence}_{(T,\tau)}$, therefore there is no additional implicit dependence on T in the bound.

Proposition 8. *For any* $\tau > 0$ *and utility function* u*, let* $\delta = \textsc{Cadence}_\tau(u)$*. Fix* $n \in \mathbb{N}_{\geq 2}$ *and define* $\mathcal{P}_T = (n, u, T, \tau, 0)$ *for all* $T > 0$*. For all* $T' \geq n\,\delta + (n-1)\tau$*, there exists* $T \in [T', T' + n(\delta+\tau)]$ *such that* $\mathrm{PoF}(\mathcal{P}_T) = 1$*.*

Proof. Fix n, τ, u, δ, and T' as in the proposition statement, and let $k = \lceil (T'+\tau)/(\delta+\tau) \rceil$. If $k \mod n = 0$, set $m = k$, otherwise (if $k \mod n > 0$), set $m = \lfloor k/n + 1 \rfloor n$. Next, set $T = m\,\delta + (m-1)\tau$. In both cases, $m \in \mathbb{R}_{\geq 1}$, $m \mod n = 0$, and $T \in [T', T' + n(\delta+\tau)]$.

Let π_T be the schedule such that $\mathcal{I}(\pi_T) = \{I_k\}_{k=1}^m$ and $I_k = ((k-1)(\delta+\tau), \delta)$ for $k = 1, \ldots, m$. The average utility of agents for π_T is

$$\frac{m\,u(\delta)}{m\delta + (m-1)\tau} = \left(\frac{m\,\delta + m\,\tau}{m\,\delta + (m-1)\tau}\right) \frac{m\,u(\delta)}{m\,\delta + m\,\tau} = \left(\frac{T+\tau}{T}\right) \frac{u(\delta)}{\delta+\tau}.$$

We show that π_T maximizes the average utility, and hence also the total utility since $T < \infty$. From Lemma 23, the optimal schedule for $\mathcal{P}_T^{UC}$ consists of intervals all with the same duration $\delta' = \textsc{Bounded-Cadence}_{(T,\tau)}(u) > 0$. Denote the number of intervals in the optimal schedule by m'. The average utility of this schedule is

$$\frac{m'\,u(\delta')}{m'\delta' + (m'-1)\tau} = \left(\frac{m'\,\delta' + m'\,\tau}{m'\,\delta' + (m'-1)\tau}\right) \frac{m'\,u(\delta')}{m'\,\delta' + m'\,\tau} = \left(\frac{T+\tau}{T}\right) \frac{u(\delta')}{\delta'+\tau}.$$

As $\arg\max_{\delta'>0} \frac{u(\delta')}{\delta'+\tau} = \delta$ by the definition of $\textsc{Cadence}_\tau(u)$, π_T is an optimal schedule for $\mathcal{P}_T^{UC}$. Since π_T consists of m intervals and $m \mod n = 0$, we can allocate the intervals of π_T in a round-robin fashion to get an envy-free schedule with the same total utility. As the envy-free schedule has the same utility as the optimal schedule, $\mathrm{PoF}(\mathcal{P}_T) = 1$. □

N.6.3. Asymptotic Recurrent Optimality in n.

Proposition 9. *Fix* u*,* $\tau > 0$*, and* T*, let* $\delta = \textsc{Bounded-Cadence}_{(T,\tau)}(u)$*, and let* $m = \frac{T+\tau}{\delta+\tau}$*. For any* $1 \leq n \leq m$*, let* $\mathcal{P}_n = (n, u, T, \tau, 0)$*. For every* $i \in [1, m]$*, there exists* $n \in \left[\left\lfloor \frac{m}{i+1} \right\rfloor, \left\lceil \frac{m}{i} \right\rceil\right]$ *such that* $\mathrm{PoF}(\mathcal{P}_n) \leq \frac{m}{m-i}$*.*

Proof. We define π^{UC} to be the schedule with m intervals all with duration δ, with total utility $m\,u(\delta)$. Note that π^{UC} is an optimal schedule for $\mathcal{P}_n^{UC}$ for any $1 \leq n \leq m$. For any $i \in [1,\, m]$, let $m = n\,i + r$ for some $n \in \left[\left\lfloor \frac{m}{i+1} \right\rfloor, \left\lceil \frac{m}{i} \right\rceil\right]$ and $0 \leq r < i$. To construct an envy-free schedule π_n^{EF} for $\mathcal{P}_n^{EF}$, we allocate $n\,i \leq m$ intervals of π^{UC} in a round-robin fashion, and leave the remaining r intervals unassigned. The total utility of π_n^{EF} is then $(m-r)\,u(\delta) \geq (m-i)\,u(\delta)$ (since $r < i$). Therefore, $\mathrm{PoF}(\mathcal{P}_n) \leq \frac{m\,u(\delta)}{(m-i)\,u(\delta)} = \frac{m}{m-i}$. □