

Dynamic Proportional Sharing: A Game-Theoretic Approach

RUPERT FREEMAN*, Duke University
SEYED MAJID ZAHEDI*, Duke University
VINCENT CONITZER, Duke University
BENJAMIN C LEE, Duke University

Sharing computational resources amortizes cost and improves utilization and efficiency. When agents pool their resources together, each becomes entitled to a portion of the shared pool. Static allocations in each round can guarantee entitlements and are strategy-proof, but efficiency suffers because allocations do not reflect variations in agents' demands for resources across rounds. Dynamic allocation mechanisms assign resources to agents across multiple rounds while guaranteeing agents their entitlements. Designing dynamic mechanisms is challenging, however, when agents are strategic and can benefit by misreporting their demands for resources.

In this paper, we show that dynamic allocation mechanisms based on max-min fail to guarantee entitlements, strategy-proofness or both. We propose the flexible lending (FL) mechanism and show that it satisfies strategy-proofness and guarantees at least half of the utility from static allocations while providing an asymptotic efficiency guarantee. Our simulations with real and synthetic data show that the performance of the flexible lending mechanism is comparable to that of state-of-the-art mechanisms, providing agents with at least 0.98x, and on average 15x, of their utility from static allocations. Finally, we propose the T -period mechanism and prove that it satisfies strategy-proofness and guarantees entitlements.

CCS Concepts: • **Information systems** → **Data centers**; • **Social and professional topics** → **Pricing and resource allocation**; • **Theory of computation** → **Algorithmic mechanism design**;

Additional Key Words and Phrases: Game theory; resource allocation; repeated game; sharing incentives; efficiency; strategy proofness

ACM Reference Format:

Rupert Freeman, Seyed Majid Zahedi, Vincent Conitzer, and Benjamin C Lee. 2018. Dynamic Proportional Sharing: A Game-Theoretic Approach. *Proc. ACM Meas. Anal. Comput. Syst.* 2, 1, Article 3 (March 2018), 36 pages. <https://doi.org/10.1145/3179406>

1 INTRODUCTION

Shared systems are defined by the competition for resources between strategic agents. In this paper, we consider a community of agents who share a non-profit system and its capital and operating costs. Sharing increases system utilization and amortizes its costs over more computation [10]. Examples include supercomputers for scientific computing [25], datacenters for Internet services [12, 32],

*These authors contributed equally to this work, ordered alphabetically.

Authors' addresses: Rupert Freeman, Duke University, rupert@cs.duke.edu; Seyed Majid Zahedi, Duke University, zahedi@cs.duke.edu; Vincent Conitzer, Duke University, conitzer@cs.duke.edu; Benjamin C Lee, Duke University, benjamin.c.lee@duke.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

2476-1249/2018/3-ART3 \$15.00

<https://doi.org/10.1145/3179406>

and clusters for academic research [8, 31]. Note, however, that our focus excludes systems in which agents explicitly pay for time on shared computational resources (i.e., infrastructure-as-a-service).

Shared systems ensure fairness by allocating resources proportionally to entitlements, which specify each agent's share of system resources relative to others [22, 24, 33]. Entitlements are dictated by exogenous factors such as agents' contributions to the shared system or priorities within the organization. A dynamic allocation mechanism should ensure agents' entitlements across time while assigning resources to computational stages that benefit most.

Guaranteeing entitlements and redistributing under-utilized resources are difficult when agents are strategic. The allocation mechanism does not know and must extract agents' utilities, which are private information. Strategic agents act selfishly to pursue their own objectives. Agents will determine whether misreporting demands can improve their performance even at the expense of others in the system. For example, an agent is likely to over-report her demand in the current time period to obtain more resources, unless doing so leads to a reduction in the resources allocated to her in later periods.

We seek allocation mechanisms that satisfy strategy-proofness (SP), which ensures that no agent benefits by misreporting her demand for resources. Strategy-proofness is a key feature contributing to efficiency as it allows the mechanism to optimize system performance according to agents' true utilities. Without SP, agents' reports may not represent their true utility and allocating based on reported demands may not produce any meaningful performance guarantee. Moreover, strategy-proof mechanisms reduce the cognitive load on agents by eliminating the need to optimally construct resource demands or preemptively respond to misreports by other agents in the system.

Strategy-proofness is complemented by sharing incentives (SI), which ensures that agents perform at least as well as they would have by not participating in the allocation mechanism (i.e., using their own resources as a smaller, private system). With sharing incentives, agents would willingly federate their resources and manage them according to the commonly agreed upon policy. A mechanism that statically enforces entitlements in every time period satisfies strategy-proofness and sharing incentives but its efficiency is poor and fails to realize the advantages of dynamic sharing across time.

In this paper, we focus on three fundamental game-theoretic desiderata: strategy-proofness, sharing incentives, and efficiency. We consider agents who derive high utility per unit of resource up until some amount of resource allocation (i.e., their demand) and derive low utility beyond that allocation. The high-low formulation is appropriate for varied resources such as processor cores, cache and memory capacity, or virtual machines in a datacenter. For example, an agent could derive high utility when additional processors permit her to dequeue more tasks from a highly critical job. Once the job's queue is empty, she derives low utility from using additional processors to replicate tasks, which guards against stragglers or failures. In another example, an agent that is allocated more power can turn on more processors, each of which provides high utility from task parallelism. Once the agent exhausts her job's parallelism, it can use additional power to boost processor voltage and frequency for lower, non-zero utility.

We propose allocation mechanisms for dynamic proportional sharing to address limitations in existing approaches. We begin by proving that policies used in state-of-the-art schedulers [2–4] fail to satisfy SP or SI. We then propose two alternative mechanisms. First, as our main contribution, we propose the flexible lending mechanism to satisfy SP, guarantee at least 50% of SI performance, and provide an asymptotic efficiency guarantee. The mechanism uses tokens to enable these theoretical guarantees. In practice, our simulations show that performance is comparable to that of state-of-the-art mechanisms and achieves 98% of SI performance, much better than the lower bound. Second, for situations where SI is a hard constraint, we propose the T-period mechanism to satisfy SP and SI while still outperforming static allocations.




Fig. 1. Users' Utility. A user derives high utility from resources up to her demand and derives low utility from resources beyond her demand.

2 PRELIMINARIES

Consider a dynamic system with n agents and R discrete rounds. Agent i contributes $e_i > 0$ units of a resource at each round, which we refer to as her *endowment*. In other words, e_i is agent i 's contribution to the federated system, which does not vary over time. Let $[n] = \{1, \dots, n\}$ and $E = \sum_{i \in [n]} e_i$ denote the total number of units to be allocated at each round. At round r , agent i has a true demand of $d_{i,r} \geq 0$ units and reports a demand of $d'_{i,r} \geq 0$. Let $\mathbf{d}'_i = (d'_{i,1}, \dots, d'_{i,R})$ denote the vector of agent i 's reports, and \mathbf{d}'_{-i} denote the reports of all agents other than i .

A dynamic allocation mechanism M assigns each agent an allocation $a_{i,r}^M(\mathbf{d}'_i, \mathbf{d}'_{-i})$ using only information from the first r entries in the demand vectors. We will often write simply $a_{i,r}$ when the exact mechanism and the demands are clear from context. Let $\mathbf{a}_i^M(\mathbf{d}'_i, \mathbf{d}'_{-i})$, often simply \mathbf{a}_i , denote the vector of agent i 's allocations. Agents have high (H) utility per resource up to their demand, and low (L) utility per resource that exceeds their demand. Formally, the utility of agent i at round r for $a_{i,r}$ units is denoted by $u_{i,r}(a_{i,r})$ and modeled as the following.

$$u_{i,r}(a_{i,r}) = \begin{cases} a_{i,r}H & \text{if } a_{i,r} \leq d_{i,r}, \\ d_{i,r}H + (a_{i,r} - d_{i,r})L & \text{if } a_{i,r} > d_{i,r}. \end{cases}$$

Figure 1 shows $u_{i,r}$ for user i with demand $d_{i,r}$ at round r . For simplicity, we assume H and L are the same for all agents, but all our results extend to the case where agents have different values of H and L (with the exception of §5.5).

While resources and demands are discrete, we allow the allocations $a_{i,r}$ to be real-valued. Real-valued allocations can be thought of as probabilistic—the realized allocation is a random allocation where agent i is allocated $a_{i,r}$ resources in expectation, which is always possible as a result of the Birkhoff-von Neumann theorem [11]. Agent i 's overall utility after R rounds for allocation \mathbf{a}_i is calculated additively as follows.

$$U_{i,R}(\mathbf{a}_i) = \sum_{r=1}^R u_{i,r}(a_{i,r}).$$

We do not consider discounting for simplicity of presentation, but our mechanisms readily extend to the case where agents discount their utilities over time.

In this paper, we focus on three main properties: strategy-proofness, sharing incentives, and efficiency. First, strategy-proofness says that agents never benefit from lying about their demands. In other words, agent i 's utility decreases if she reports $\mathbf{d}'_i \neq \mathbf{d}_i$.

DEFINITION 1. *Mechanism M satisfies strategy-proofness (SP) if*

$$U_{i,R}(\mathbf{a}_i^M(\mathbf{d}_i, \mathbf{d}'_{-i})) \geq U_{i,R}(\mathbf{a}_i^M(\mathbf{d}'_i, \mathbf{d}'_{-i})) \quad \forall i, \forall R, \forall \mathbf{d}_i, \forall \mathbf{d}'_i, \text{ and } \forall \mathbf{d}'_{-i}.$$

Next, sharing incentives says that by participating in the mechanism, agents receive at least the utility they would have received by not participating.

DEFINITION 2. *Mechanism M satisfies sharing incentives (SI) if*

$$U_{i,R}(a_i^M(\mathbf{d}_i, \mathbf{d}'_{-i})) \geq U_{i,R}(e_i) \quad \forall i, \forall R, \forall \mathbf{d}_i, \text{ and } \forall \mathbf{d}'_{-i}.$$

We also define a relaxed notion of α -sharing incentives, which says that every agent gets at least an α fraction of the utility that she would have received without taking part in the mechanism. Note that 1-SI is equivalent to SI.

DEFINITION 3. *Mechanism M satisfies α -SI if*

$$U_{i,R}(a_i^M(\mathbf{d}_i, \mathbf{d}'_{-i})) \geq \alpha U_{i,R}(e_i) \quad \forall i, \forall R, \forall \mathbf{d}_i, \text{ and } \forall \mathbf{d}'_{-i}.$$

Finally, efficiency says that all resources should be allocated, and an agent with L valuation should never receive a resource while there are agents with H valuation for that resource.

DEFINITION 4. *Mechanism M satisfies efficiency if*

$$\sum_{i \in [n]} a_{i,r}^M = E,$$

and if $a_{i,r}^M > d'_{i,r}$ for some agent i and round r , then $a_{j,r}^M \geq d'_{j,r}$ for all agents.

Note that efficiency is relative to the agents' reports, not their actual valuations, which are hidden from the mechanism. Therefore, in situations where agents lie about their valuations, it is possible that even an efficient mechanism allocates a unit inefficiently with respect to the actual valuations. With this in mind, there is little value in a mechanism that is efficient but not SP. Similarly, if a mechanism does not satisfy SI, then agents may not want to participate in it. So an efficient mechanism that does not satisfy SI may not actually exhibit efficiency gains in practice because agents choose not to participate. In some contexts, SI may not be of concern because agents are forced to participate or are willing to risk participation if gains are likely large and losses are likely small.

For readability, some proofs are omitted and appear in the appendix.

3 EXISTING MECHANISMS

In this section, we focus on the (weighted) max-min fairness policy, which is one of the most widely used policies in computing systems. It is deployed in many state-of-the-art datacenter schedulers such as the Hadoop Fair Scheduler [3], Hadoop Capacity Scheduler [2] and Spark Dynamic Allocator [4]. And it has been extensively studied in the literature [16, 17, 29].

A dynamic allocation mechanism could deploy the max-min policy for two different objectives: maximizing the minimum accumulated allocations up to a round, or maximizing the minimum allocation at each round, independently of previous rounds. We call the first mechanism *Dynamic Max-Min (DMM)* and the second mechanism *Static Max-Min (SMM)*. First, at each round r , DMM selects the allocation that maximizes $\min_i \sum_{r'=1}^r a_{i,r'}/e_i$, the minimum weighted cumulative allocation; subject to this, it maximizes the second lowest weighted cumulative allocation, and so on. This maximization is subject to the constraint that no resource is allocated to an agent with low valuation as long as there are agents with high valuation.

Second, at each round r , SMM selects the allocation that maximizes $\min_i a_{i,r}/e_i$, the minimum weighted allocation at that round; subject to this, it maximizes the second lowest weighted allocation, and so on. This maximization is also subject to the constraint that no resource is allocated to an agent with low valuation as long as there are agents with high valuation. Under SMM, agents are

guaranteed to receive their demands as long as they are less than or equal to their endowment. Agents with demands higher than their endowments receive extra resources from agents with demands lower than their endowments. Unlike DMM, SMM allocates resources locally at round r , regardless of agents' allocations prior to round r .

In the rest of this section, we study properties of these two mechanisms. In particular, we focus on three properties: strategy-proofness, sharing incentives, and efficiency. We examine whether the existing mechanisms satisfy these properties for the special case when $L = 0$ and for the general case when $L > 0$.

3.1 Properties of Mechanisms for $L = 0$

When $L = 0$, one might think that agents do not have any incentive to misreport their demands. However, we show that DMM fails to satisfy SI and SP.

THEOREM 5. *Dynamic max-min mechanism violates sharing incentives, even when $L = 0$.*

PROOF. Suppose that $R = 10$ and there are three agents, each with $e_i = 3$. For all rounds $r \neq 10$, the demands are $d_{1,r} = 1$, $d_{2,r} = 2$, and $d_{3,r} = 6$. For rounds $r = 1, \dots, 9$, each agent is allocated exactly her demand. After round 9, utilities for agents 1, 2 and 3 are $9H$, $18H$ and $54H$, respectively. At round 10, demands are $d_{1,10} = 9$, $d_{2,10} = 9$, and $d_{3,10} = 6$. DMM allocates all 9 units to agent 1, which maximizes the minimum weighted cumulative allocation. Consider agent 2. Under DMM, agent 2's allocation is $a_{2,r} = 2$ for all $r \neq 10$ and $a_{2,10} = 0$. If she had not participated in the mechanism, then she would have obtained the same utility in each round $r \neq 10$, but a strictly higher utility in round $r = 10$. \square

THEOREM 6. *Dynamic max-min mechanism violates strategy-proofness, even when $L = 0$ [6].*

PROOF. Consider three agents with equal endowments $m_1 = m_2 = m_3 = 1$ sharing three units of a resource for three rounds. The demand of agent 1 is 3 for all three rounds. Agent 2's demand is 3 for rounds 1 and 3 and 0 for round 2. And agent 3 has a demand of 3 for round 2 and 0 for rounds 1 and 3. Agent 1 achieves utility of $3.375H$ by truthful reporting. If agent 1 misreports 0 for round 1, her utility would increase to $3.75H$. \square

Since DMM does not satisfy SP, it cannot guarantee any meaningful notion of efficiency, as explained in §2. Next, we show that SMM satisfies SI, SP, and efficiency.

THEOREM 7. *Static max-min mechanism satisfies strategy-proofness, sharing incentives, and efficiency when $L = 0$.*

PROOF. We start by proving that SMM satisfies SP. Under SMM, allocations at round r are independent of allocations at previous rounds. Suppose that agent i reports $d'_{i,r} \neq d_{i,r}$ at round r . Let $a'_{i,r}$ and $a_{i,r}$ denote i 's allocations at round r for reporting $d'_{i,r}$ and $d_{i,r}$, respectively. If $a_{i,r} \geq d_{i,r}$, then i already receives her highest possible utility, $d_{i,r}H$ (because $L = 0$), and she cannot benefit from misreporting.

If $a_{i,r} < d_{i,r}$, then for all $j \neq i$, we have: (1) $a_{j,r} \leq d_{j,r}$ and (2) $a_{i,r}/e_i \geq a_{j,r}/e_j$. The former holds by SMM's definition. The latter holds because SMM maximizes the minimum weighted allocations in a lexicographical order. If there is j with $a_{j,r}/e_j > a_{i,r}/e_i$, then SMM should decrease $a_{j,r}$ and increase $a_{i,r}$. Now, suppose for contradiction that $a'_{i,r} > a_{i,r}$. Since $\sum_k a'_{k,r} = \sum_k a_{k,r}$, there should be an agent ℓ with $a'_{\ell,r} < a_{\ell,r} \leq d_{\ell,r}$. Therefore, we have:

$$a'_{\ell,r}/e_\ell < a_{\ell,r}/e_\ell \leq a_{i,r}/e_i < a'_{i,r}/e_i.$$

This is a contradiction because SMM could improve its objective value by decreasing $a'_{i,r}$ and increasing $a'_{\ell,r}$.

To see that SMM satisfies SI, note that an agent can guarantee herself at least e_i resources (her utility from not participating) at each round by reporting $d'_{i,r} = e_i$ for all r . By SP, truthful reporting achieves at least this utility. Therefore, truthful reporting achieves at least as much utility as not participating in SMM, which proves SI. Finally, SMM satisfies efficiency by definition, since it either completely fulfills all demands or allocates all resources to agents that value them highly. \square

3.2 Properties of Mechanisms for $L > 0$

We now consider the general setting where an agent's low valuation is still positive. Unfortunately, SMM no longer retains its properties from the $L = 0$ case. Agents are no longer indifferent to forsaking low-valued resources and may lie in order to receive them.

THEOREM 8. *When $L > 0$, static max-min mechanism violates strategy-proofness and sharing incentives.*

PROOF. Consider an instance with 2 agents, each with endowment $e_i = 1$, and a single round. Agent 1 has demand 2 and agent 2 has demand 0. SMM allocates both resources to agent 1 and nothing to agent 2. However, had agent 2 not participated in the mechanism, she would have received one resource and utility $L > 0$. Similarly, had she misreported her demand to be 1, she would have received one resource and utility $L > 0$. \square

Indeed, in this general setting, no mechanism can simultaneously satisfy efficiency and either of the two other desired properties.

THEOREM 9. *When $L > 0$, there is no dynamic mechanism that satisfies α -sharing incentives and efficiency, for any $\alpha > 0$.*

PROOF. Consider an instance with two agents, each with endowment $e_i = 1$, and a single round. Agent 1 has demand 2 and agent 2 has demand 0. Efficiency dictates that we allocate both resources to agent 1, which would violate α -SI for agent 2 for any $\alpha > 0$. \square

THEOREM 10. *When $L > 0$, there is no dynamic mechanism that satisfies strategy-proofness and efficiency.*

PROOF. Consider an instance with two agents, each with endowment $e_i = 1$, and a single round. Both agents have demand 0. For efficiency, the mechanism must allocate all the resources so that at least one agent receives $a_{i,1} > 0$. Supposing without loss of generality that $a_{1,1} > 0$, then $a_{2,1} < 2$. If agent 2 misreports $d'_{2,1} = 2$, by efficiency, the mechanism must allocate both resources to agent 2, which is an improvement over her utility from reporting truthfully. \square

Note that SP and SI are compatible. A mechanism that statically allocates agents their endowments satisfies SP and SI; agents have no incentive to misreport because allocations do not depend on reports and agents receive their fair share of resources. This mechanism clearly fails to satisfy efficiency and does not extract any benefit from sharing. In §5, we propose a mechanism that satisfies strategy-proofness, guarantees each user at least 50% of their utilities from sharing incentives, and provides an asymptotic efficiency guarantee.

4 PROPORTIONAL SHARING WITH CONSTRAINTS PROCEDURE

The mechanisms we present in the remainder of this paper have, at their core, a procedure we call *Proportional Sharing With Constraints (PSWC)*. The procedure allocates some amount of resources among agents proportionally to their (exogenous) weights subject to (agent-dependent) minimum




Fig. 2. Proportional Sharing With Constraints. There are six agents with equal weights. The allocations are represented by the height of the corresponding blue vertical bar. The allocations can be thought of as the ‘most equal’ allocations, subject to no agent receiving less than her minimum constraint or more than her maximum.

and limit constraints: (1) each agent receives at least her minimum allocation, and (2) each agent should receives no more than her limit allocation.

Formally, PSWC takes as input an amount to allocate A , weights $\mathbf{w} = (w_1, \dots, w_n)$, minimum allocations $\mathbf{m} = (m_1, \dots, m_n)$, and limit allocations $\mathbf{l} = (l_1, \dots, l_n)$. PSWC outputs a vector of allocations $\mathbf{a} = (a_1, \dots, a_n)$ defined as the solution to the following program.

$$\begin{aligned}
 & \text{Minimize } x, \\
 & \text{s.t. } a_i/w_i \leq x && \text{if } m_i < a_i \leq l_i, \\
 & a_i \leq l_i && \forall i, \\
 & a_i \geq m_i && \forall i, \\
 & \sum_{i \in [n]} a_i = A.
 \end{aligned}$$

PSWC is illustrated in Figure 2. The program can be solved in $O(n \log(n))$ time by the Divvy algorithm [20]. The Divvy algorithm proceeds by sorting the limit and minimum allocation bounds in $O(n \log(n))$ time, and then conducting a linear time search for the optimal value of x by increasing the allocations in discrete steps until all resources have been allocated.

The following lemma characterizes the allocations produced by the PSWC procedure and will be useful in our later proofs.

LEMMA 11. *Under PSWC, for every agent i , $a_i = \max(m_i, \min(l_i, xw_i))$.*

PROOF. First, we show that if $m_i < xw_i$, then $a_i = \min(l_i, xw_i)$. If $a_i > \min(l_i, xw_i)$, then at least one constraint is violated. If $a_i < \min(l_i, xw_i)$, then there exists at least one agent ℓ such that $a_\ell = xw_\ell$ because otherwise, x is not optimal. In this case, a_i can be increased while a_ℓ for all ℓ with $a_\ell = xw_\ell$ decreases. This allows for a smaller value of x , which contradicts the optimality of x . Next, we show that if $m_i \geq xw_i$, then $a_i = m_i$. Since a_i cannot be less than m_i , if a_i is not equal to m_i , then $a_i > m_i$, which means $a_i > xw_i$. However, since $a_i > m_i$, the first constraint dictates that $a_i \leq xw_i$, a contradiction. Combining these two cases gives the desired result. \square

Our proposed mechanisms all have similar structure. First, agents always receive exactly the same number of resources that they contribute to the system (over the entire R rounds). This is a fairness primitive in its own right, but is primarily a design feature that helps us provide desirable properties. Second, all our proposed dynamic mechanisms call the PSWC procedure to allocate

resources at each round. Our mechanisms are determined primarily by how we set the minimum and maximum constraints.

5 FLEXIBLE LENDING MECHANISM

We now turn to designing mechanisms that satisfy our game-theoretic desiderata while increasing efficiency significantly over static allocation. The static allocation mechanism satisfies both SP and SI, but it does not exhibit any gains from sharing. DMM and SMM sacrifice SP and SI in exchange for efficiency. However, in the absence of SP, any guarantee on efficiency based on agents' demands is not meaningful as agents have incentives to misreport their demands when $L > 0$. In this section, we present the flexible lending (FL) mechanism. The flexible lending mechanism achieves strategy-proofness and an asymptotic efficiency guarantee. FL satisfies a theoretical 0.5 approximation to SI and our simulation results show that it significantly outperforms this bound in practice (see §7).

5.1 Definition

For a fixed number of rounds R , FL allocates exactly Re_i resources to each agent i , which is exactly her contribution to the shared pool over all R rounds. The mechanism enforces this constraint by simply removing agent i from the list of eligible agents once she receives Re_i resources in total. We keep track of the resources each agent has received with a running token count t_i , effectively 'charging' each agent a token for every resource she receives. We denote by $t_{i,r}$ the number of tokens that agent i holds at the start of round r . Thus, the number of tokens that an agent holds puts a hard limit on the number of resources she can receive at any given round.

Algorithm 1 presents the flexible lending mechanism. We define \bar{d}_i to be the allocatable demand of agent i at each round, which is simply the minimum of her reported demand $d'_{i,r}$ and the number of tokens she has remaining t_i . We distinguish between two cases depending on whether the total allocatable demand is higher or lower than the total supply of resources.

First, if the total allocatable demand is at least as high as the total supply, then FL runs PSWC with the minimum allocation for each agent set to 0, and the limit allocation set to \bar{d}_i . This way, resources are allocated proportionally among all agents that want them. Second, if the total allocatable demand is less than the total supply, then agents receive their full allocatable demand. Therefore, FL runs PSWC with minimum allocation for each agent i set to \bar{d}_i , and limit allocations set to her number of tokens t_i (which is always at least as large as her allocatable demand). This way, FL allocates resources proportionally among all agents, subject to the condition that no agent receives fewer resources than her demand.

Algorithm 1 Flexible Lending Mechanism

```

t =  $Re$  ▷ Initialize token count
for  $r \in \{1, \dots, R\}$  do
   $\bar{\mathbf{d}} \leftarrow \min(\mathbf{d}'_{\cdot,r}, \mathbf{t})$  ▷  $\bar{d}_i$  is  $i$ 's allocatable demand
   $D \leftarrow \sum_{i \in [n]} \bar{d}_i$ 
  if  $D \geq E$  then
     $\mathbf{a}_{\cdot,r} \leftarrow \text{PSWC}(A = E, \mathbf{l} = \bar{\mathbf{d}}, \mathbf{m} = \mathbf{0}, \mathbf{w} = \mathbf{e})$ 
  else
     $\mathbf{a}_{\cdot,r} \leftarrow \text{PSWC}(A = E, \mathbf{l} = \mathbf{t}, \mathbf{m} = \bar{\mathbf{d}}, \mathbf{w} = \mathbf{e})$ 
  end if
   $\mathbf{t} \leftarrow \mathbf{t} - \mathbf{a}_{\cdot,r}$ 
end for

```

We illustrate FL with an example.

EXAMPLE 12. Consider a system with three agents and four rounds. Each agent has endowment $e_i = 1$. Suppose that agents' (truthful) reports are given by the following table:

	$d_{i,1}$	$d_{i,2}$	$d_{i,3}$	$d_{i,4}$
$i = 1$	3	1	1	0
$i = 2$	0	2	1	2
$i = 3$	0	0	0	4

FL allocations are given by the following table:

	$a_{i,1}^{FL}$	$a_{i,2}^{FL}$	$a_{i,3}^{FL}$	$a_{i,4}^{FL}$
$i = 1$	3	1	0	0
$i = 2$	0	2	1.5	0.5
$i = 3$	0	0	1.5	2.5

While all agents have tokens remaining, FL efficiently allocates resources. However, in round 3, agent 1 has no tokens remaining and therefore the supply of resources exceeds the allocatable demand. In this case, resources are evenly divided between agents 2 and 3. In the final round, agent 2 can receive only 0.5 resources before running out of tokens, so the rest of the resources are allocated to agent 3.

5.2 Basic Properties

Next, we study the properties of FL. We first show that FL satisfies strategy-proofness. We then show that FL guarantees at least 50% of SI performance. And finally we show that FL provides an asymptotic efficiency guarantee. Throughout this section, we extensively use the following lemma which characterizes FL allocations.

LEMMA 13. Let x denote the objective value of FL's call to PSWC at round r . If $D \geq E$, then $a_{i,r} = \min(xe_i, d_{i,r}, t_{i,r})$. If $D < E$, then $a_{i,r} = \min(t_{i,r}, \max(d_{i,r}, xe_i))$.

PROOF. Suppose first that $D \geq E$. Substituting the relevant terms into Lemma 11, we have

$$a_{i,r} = \max(0, \min(\min(d_{i,r}, t_{i,r}), xe_i)) = \min(xe_i, d_{i,r}, t_{i,r}).$$

If instead $D < E$, then again substituting into Lemma 11 gives

$$a_{i,r} = \max(\min(d_{i,r}, t_{i,r}), \min(t_{i,r}, xe_i)) = \min(t_{i,r}, \max(d_{i,r}, xe_i)).$$

The final equality, $\max(\min(A, B), \min(A, C)) = \min(A, \max(B, C))$ can easily be checked to hold case by case for any relative ordering of A , B , and C . \square

We next prove a basic monotonicity result, which states that if we shift some tokens to a single agent from all other agents, then the agent with more tokens achieves a (weakly) higher allocation. The proof follows easily from Lemma 13 and is deferred to the Appendix.

LEMMA 14. Consider some agent i , and suppose that $t'_{i,r} \geq t_{i,r}$, $t'_{j,r} \leq t_{j,r}$ for all $j \neq i$, and $d'_{k,r} = d_{k,r}$ for all $k \in [n]$. Then $a'_{i,r} \geq a_{i,r}$.

As our main technical result, we show in the following subsection that FL is strategy-proof. At a high level, we show that if an agent receives fewer high-valued resources as a result of misreporting, then her allocations in all future rounds are weakly higher. This means that she cannot receive fewer low-valued resources at any future round, relative to her allocations had she not misreported. Therefore, because the total number of resources allocated to each agent is fixed (by the initial token count), her misreport can only result in trading high-valued resources at an early round for other, potentially low-valued, resources at later rounds.

5.3 Strategy-Proofness

Suppose agent i reports demands that are not equal to her true demands. Let r' be the latest round for which i misreports. That is, $r' = \max\{r : d'_{i,r} \neq d_{i,r}\}$. Suppose that $d'_{i,r'} < d_{i,r'}$. We show that, all else being equal, i could (weakly) improve her utility by instead reporting $d'_{i,r'} = d_{i,r'}$. The proof that reporting $d'_{i,r'} > d_{i,r'}$ is also (weakly) worse than reporting $d'_{i,r'} = d_{i,r'}$ is almost identical and can be found in Appendix B. It follows from this that FL is strategy-proof, since any non-truthful reports can be converted to truthful reports one round at a time, (weakly) improving i 's utility.

We consider parallel universes: one in which agent i misreports $d'_{i,r'}$ at round r' (the 'misreported instance') and one in which she truthfully reports $d_{i,r}$ (the 'truthful instance,' even though i 's reports prior to r' may yet be non-truthful). All other reports are identical in both universes. We denote allocations and tokens in the misreported instance using a' and t' , respectively, and in the truthful instance by a and t . We denote by D_r and D'_r the total demand D at round r in the truthful and misreported instances, respectively.

We first note that for all rounds prior to r' , the allocations in the truthful and misreported instances are the same.

LEMMA 15. *For all rounds $r < r'$ and for all agents j , $a'_{j,r} = a_{j,r}$.*

PROOF. The mechanism does not take future reports into account, so because agents' demands in both instances are identical up to round r' , so are the allocations. \square

We next show a monotonicity lemma, which says that agent i 's allocation at round r' is (weakly) smaller in the misreported instance than the truthful instance, and all other agents' allocations are (weakly) larger.

LEMMA 16. *For all agents $j \neq i$, we have that $a'_{j,r'} \geq a_{j,r'}$. Further, $a'_{i,r'} \leq a_{i,r'}$.*

PROOF. We prove the statement for all $j \neq i$. The statement for i follows immediately because the total number of allocated resources is fixed. Observe first that

$$D'_{r'} = \sum_{k \in [n]} \min(d'_{k,r'}, t_{k,r'}) \leq \sum_{k \in [n]} \min(d_{k,r'}, t_{k,r'}) = D_{r'}$$

since i 's demand decreases in the misreported instances but all other demands and token counts stay the same. Let x' denote the objective value in FL's call to PSWC in the misreported instance, and x in the truthful instance.

Suppose that $E \leq D'_{r'} \leq D_{r'}$. Suppose first that $x' > x$. Then, by Lemma 13, for all $j \neq i$, we have

$$a'_{j,r'} = \min(x'e_j, d_{j,r'}, t_{j,r'}) \geq \min(xe_j, d_{j,r'}, t_{j,r'}) = a_{j,r'}$$

Next, suppose that $x' \leq x$. Then, again by Lemma 13 and the fact that $d'_{i,r'} < d_{i,r'}$,

$$a'_{i,r'} = \min(x'e_i, d'_{i,r'}, t_{i,r'}) \leq \min(xe_i, d_{i,r'}, t_{i,r'}) = a_{i,r'}$$

And, for all $j \neq i$,

$$a'_{j,r'} = \min(x'e_j, d_{j,r'}, t_{j,r'}) \leq \min(xe_j, d_{j,r'}, t_{j,r'}) = a_{j,r'}$$

Because $a'_{k,r'} \leq a_{k,r'}$ for all agents k , and $\sum_{k \in [n]} a_{k,r'} = \sum_{k \in [n]} a'_{k,r'}$, it must be the case that $a'_{k,r'} = a_{k,r'}$ for all k , which satisfies the statement of the lemma.

Next, suppose that $D'_{r'} < E \leq D_{r'}$. By the definition of FL, $a'_{k,r'} \geq \min(d'_{k,r'}, t_{k,r'})$ for all k , and $a_{k,r'} \leq \min(d_{k,r'}, t_{k,r'})$ for all k . Since $\min(d'_{j,r'}, t_{j,r'}) = \min(d_{j,r'}, t_{j,r'})$ for all $j \neq i$, we have that $a'_{j,r'} \geq a_{j,r'}$, implying also that $a'_{i,r'} \leq a_{i,r'}$.

Finally, suppose that $D'_{r'} \leq D_{r'} < E$. Suppose first that $x' \leq x$. Then, by Lemma 13 and the assumption that $d'_{i,r'} < d_{i,r'}$, we have

$$a'_{i,r'} = \min(t_{i,r'}, \max(x'e_i, d'_{i,r'})) \leq \min(t_{i,r'}, \max(xe_i, d_{i,r'})) = a_{i,r'}$$

and

$$a'_{j,r'} = \min(t_{j,r'}, \max(x'e_j, d'_{j,r'})) \leq \min(t_{j,r'}, \max(xe_j, d_{j,r'})) = a_{j,r'}$$

for all $j \neq i$. Because $a'_{k,r'} \leq a_{k,r'}$ for all agents k , and $\sum_{k \in [n]} a_{k,r'} = \sum_{k \in [n]} a'_{k,r'}$, it must be the case that $a'_{k,r'} = a_{k,r'}$ for all k , which satisfies the lemma's statement. Next, suppose that $x' > x$. Then, again by Lemma 13, for all $j \neq i$, we have

$$a'_{j,r'} = \min(t_{j,r'}, \max(x'e_j, d_{j,r'})) \geq \min(t_{j,r'}, \max(xe_j, d_{j,r'})) = a_{j,r'}. \quad \square$$

If it is the case that $a'_{i,r'} = a_{i,r'}$, then it must also be the case that $a'_{j,r'} = a_{j,r'}$ for all $j \neq i$. That is, allocations at round r' are the same in the misreported instance as the truthful instance. Therefore, for all rounds $r \leq r'$, allocations in both universes would be the same. In all rounds $r > r'$, reports in both universes are the same. Together, these imply that allocations for all rounds $r > r'$ would be the same in both universes. In particular, i does not profit from her misreport and could weakly improve her utility by reporting $d'_{i,r'} = d_{i,r'}$. So, for the remainder of this section, we assume that $a'_{i,r'} < a_{i,r'}$.

Our next lemma states that the resources that i sacrifices in round r' are *high*-valued resources for her. The intuition is that if it were the case that i was being forced to receive low-valued resources under truthful reporting, then she will still be forced to receive the same number of resources when she under-reports her demand (since there is no agent with excess demand to absorb extra resources).

LEMMA 17. *If $a'_{i,r'} < a_{i,r'}$, then $a_{i,r'} \leq d_{i,r'}$.*

PROOF. Suppose for contradiction that $a_{i,r'} > d_{i,r'}$. It must therefore be the case that $D'_{r'} \leq D_{r'} < E$, where the first inequality holds because $d'_{j,r'} = d_{j,r'}$ for all $j \neq i$ and $d'_{i,r'} < d_{i,r'}$. Let x denote the objective value of FL's call to PSWC in the truthful instance, and x' in the misreported instance. Suppose that $x' \leq x$. Then, by Lemma 13 and the assumption that $d'_{i,r'} < d_{i,r'}$,

$$a'_{i,r'} = \min(t_{i,r'}, \max(x'e_i, d'_{i,r'})) \leq \min(t_{i,r'}, \max(xe_i, d_{i,r'})) = a_{i,r'},$$

and for all $j \neq i$,

$$a'_{j,r'} = \min(t_{j,r'}, \max(x'e_j, d_{j,r'})) \leq \min(t_{j,r'}, \max(xe_j, d_{j,r'})) = a_{j,r'}.$$

Because $a'_{k,r'} \leq a_{k,r'}$ for all agents k , and $\sum_{k \in [n]} a_{k,r'} = \sum_{k \in [n]} a'_{k,r'}$, it must be the case that $a'_{k,r'} = a_{k,r'}$ for all k . This contradicts the assumption that $a'_{i,r'} < a_{i,r'}$.

Now suppose that $x' > x$. Note that $xe_i > d_{i,r'} > d'_{i,r'}$, where the first inequality holds because $a_{i,r'} > d_{i,r'}$. Then, again by Lemma 13 and the previous observation, we have

$$\begin{aligned} a'_{i,r'} &= \min(t_{i,r'}, \max(x'e_i, d'_{i,r'})) = \min(t_{i,r'}, x'e_i) \\ &\geq \min(t_{i,r'}, xe_i) = \min(t_{i,r'}, \max(xe_i, d_{i,r'})) = a_{i,r'}, \end{aligned}$$

which contradicts $a'_{i,r'} < a_{i,r'}$. Since we arrive at a contradiction in all cases, the lemma statement must be true. \square

As a corollary, we can write the difference in utility between the truthful and misreported instances that i derives from round r' .

COROLLARY 18. $u_{i,r'}(a_{i,r'}) - u_{i,r'}(a'_{i,r'}) = H(a_{i,r'} - a'_{i,r'})$.

PROOF. Because $a'_{i,r'} < a_{i,r'} \leq d_{i,r'}$, we can substitute the utility values from Equation (2):

$$u_{i,r'}(a_{i,r'}) - u_{i,r'}(a'_{i,r'}) = a_{i,r'}H - a'_{i,r'}H = H(a_{i,r'} - a'_{i,r'}).$$

□

For a fixed agent k , denote by r_k the round at which agent k runs out of tokens in the truthful instance. That is, r_k is the first (and only) round with $a_{r_k} = t_{k,r_k} > 0$. Note that $r_i \geq r'$, since $a_{i,r'} > 0$. Given this, our next lemma states that, under certain conditions, the effect of i 's misreport, $d'_{i,r} < d_{i,r}$, is to increase the objective value of FL's call to PSWC.

LEMMA 19. *Let $r < r_i$ (i.e. $a_{i,r} < t_{i,r}$). Suppose $t'_{j,r} \leq t_{j,r}$ for all agents $j \neq i$. Suppose that either $\min(D_r, D'_r) \geq E$ or $\max(D_r, D'_r) < E$. Then $x' \geq x$, where x' denotes the objective value of FL's call to PSWC in the misreported instance and x in the truthful instance.*

PROOF. First, suppose that $\min(D_r, D'_r) \geq E$. Suppose for contradiction that $x' < x$. By Lemma 13, for all $j \neq i$,

$$a'_{j,r} = \min(x'e_j, d_{j,r}, t'_{j,r}) \leq \min(xe_j, d_{j,r}, t_{j,r}) = a_{j,r},$$

where the inequality follows from the assumption that $x' < x$ and that $t'_{j,r} \leq t_{j,r}$. Further,

$$a'_{i,r} = \min(x'e_i, d_{i,r}, t'_{i,r}) \leq \min(x'e_i, d_{i,r}) \leq \min(xe_i, d_{i,r}) = \min(xe_i, d_{i,r}, t_{i,r}) = a_{i,r},$$

where the second inequality follows from the assumption that $x' < x$, and the second to the last equality follows from the assumption that $a_{i,r} < t_{i,r}$. Therefore, $a'_{k,r} \leq a_{k,r}$ for all agents k . Since $\sum a'_{k,r} = \sum a_{k,r}$, it must be the case that $a'_{k,r} = a_{k,r}$ for all agents k . Now, by the definition of FL in this case, $a_{k,r}/e_k \leq x' < x$ for all agents k with $a_{k,r} > 0$. Therefore x is not the optimal objective value of PSWC in the truthful instance, a contradiction. Thus, $x' \geq x$.

Next, suppose that $\max(D_r, D'_r) < E$. Suppose for contradiction that $x' < x$. By Lemma 13,

$$a'_{j,r} = \min(t'_{j,r}, \max(x'e_j, d_{j,r})) \leq \min(t_{j,r}, \max(xe_j, d_{j,r})) = a_{j,r},$$

for all $j \neq i$, where the inequality follows from the assumption that $x' < x$ and that $t'_{j,r} \leq t_{j,r}$. Further, we have

$$\begin{aligned} a'_{i,r} &= \min(t'_{i,r}, \max(x'e_i, d_{i,r})) \leq \max(x'e_i, d_{i,r}) \\ &\leq \max(xe_i, d_{i,r}) = \min(t_{i,r}, \max(xe_i, d_{i,r})) = a_{i,r}, \end{aligned}$$

where the second inequality follows from the assumption that $x' < x$ and the second to last equality from the assumption $a_{i,r} < t_{i,r}$. Therefore, $a'_{k,r} \leq a_{k,r}$ for all agents k . Since $\sum a'_{k,r} = \sum a_{k,r}$, it must be the case that $a'_{k,r} = a_{k,r}$ for all agents k . Consider all agents with $\min(d_{k,r}, t_{k,r}) < a_{k,r}$ (i.e. those agents for which the first constraint in the PSWC program binds in the truthful instance). For all such agents, we have

$$\min(d_{k,r}, t_{k,r}) < a_{k,r} \implies d_{k,r} < a_{k,r} \leq t_{k,r} \implies d_{k,r} < a'_{k,r} \leq t'_{k,r} \implies \min(d_{k,r}, t'_{k,r}) < a'_{k,r},$$

which implies that the constraints bind in the misreported instance as well. Therefore, $a'_{k,r}/e_k \leq x' < x$ for all agents k for which the first constraint binds in the truthful instance. Therefore x is not the optimal objective value of the PSWC program in the truthful instance, a contradiction. Thus, $x' \geq x$. □

Using Lemma 19, we show our main lemma. This lemma allows us to make an inductive argument that, after giving up some resources in round r' , i 's allocation is (weakly) larger for all future rounds in the misreported instance than the truthful instance.

LEMMA 20. Let $r' < r < r_i$ (i.e. $a_{i,r} < t_{i,r}$). Suppose that $t'_{j,r} \leq t_{j,r}$ for all agents $j \neq i$. Then for all $j \neq i$, either: (1) $a'_{j,r} = t'_{j,r}$, or (2) $a'_{j,r} \geq a_{j,r}$.

PROOF. Note that $t'_{j,r} \leq t_{j,r}$ for all $j \neq i$ implies that $t'_{i,r} \geq t_{i,r}$, which we use in the proof. Also, because $r > r'$, we know that $d'_{i,r} = d_{i,r}$, as r' is the last round for which $d'_{i,r} \neq d_{i,r}$. We assume that condition (1) from the lemma statement is false (i.e. $a'_{j,r} < t'_{j,r}$) and show that condition (2) must hold. Suppose first that $D_r < E$. Then, because $a_{i,r} < t_{i,r}$, we know that $d_{i,r} \leq t_{i,r} \leq t'_{i,r}$. This implies that $\min(d_{i,r}, t_{i,r}) = \min(d_{i,r}, t'_{i,r}) = d_{i,r}$. Let $j \neq i$. Since $t'_{j,r} \leq t_{j,r}$, we have $\min(d_{j,r}, t'_{j,r}) \leq \min(d_{j,r}, t_{j,r})$. Therefore, it is the case that $D'_r \leq D_r < E$. By Lemma 13 and the assumption that $a'_{j,r} < t'_{j,r}$, it must be the case that $a'_{j,r} = \max(d_{j,r}, x'e_j)$. Further, by Lemma 19, we know that $x' \geq x$. Therefore, we have

$$a_{j,r} = \max(d_{j,r}, xe_j) \leq \max(d_{j,r}, x'e_j) = a'_{j,r}.$$

That is, condition (2) from the lemma statement holds.

Now suppose that $D_r \geq E$. Then, from the definition of the mechanism, we have that $a_{j,r} \leq \min(d_{j,r}, t_{j,r}) \leq d_{j,r}$. If it is the case that $D'_r < E$, then we have that $a'_{j,r} \geq \min(d_{j,r}, t'_{j,r}) = d_{j,r}$, where the equality holds because otherwise we would have $a'_{j,r} \geq t'_{j,r}$, violating the assumption that $a'_{j,r} < t'_{j,r}$. Using these inequalities, we have $a'_{j,r} \geq d_{j,r} \geq a_{j,r}$, so condition (2) from the statement of the lemma holds. Finally, it may be the case that $D_r \geq E$ and $D'_r \geq E$. By Lemma 13 and the assumption that $a'_{j,r} < t'_{j,r}$, we have

$$a'_{j,r} = \min(d_{j,r}, x'e_k) \geq \min(d_{j,r}, xe_k) = a_{j,r},$$

where the inequality follows from Lemma 19. Thus, condition (2) of the lemma statement holds. \square

Finally, we prove that the flexible lending mechanism is strategy-proof. This proof establishes that misreporting $d'_{i,r}$ is never beneficial for an agent.

THEOREM 21. *The flexible lending mechanism satisfies SP.*

PROOF. We first observe that for every $r \leq r_i$, $t'_{j,r} \leq t_{j,r}$ for every $j \neq i$. This is true for every $r \leq r'$ because $a'_{j,r} = a_{j,r}$ for $r < r'$, by Lemma 15. For $r = r' + 1$, it follows from Lemma 16, which says that $a'_{j,r'} \geq a_{j,r'}$. For all subsequent rounds, up to and including $r = r_i$, it follows inductively from Lemma 20: $t'_{j,r} \leq t_{j,r}$ implies that either $a'_{j,r} = t'_{j,r}$, in which case $t'_{j,r+1} = 0 \leq t_{j,r+1}$, or $a'_{j,r} \geq a_{j,r}$, in which case $t'_{j,r+1} = t'_{j,r} - a'_{j,r} \leq t_{j,r} - a_{j,r} = t_{j,r+1}$.

Consider an arbitrary round $r \neq r'$, with $r \leq r_i$. By the above argument, we know that $t'_{j,r} \leq t_{j,r}$ for all $j \neq i$. Further, because reports in the truthful and misreported instances are identical on all rounds $r \neq r'$, we have that $d_{k,r} = d'_{k,r}$ for all $k \in [n]$. Therefore, by Lemma 14, $a'_{i,r} \geq a_{i,r}$. For rounds $r > r_i$, it is also true that $a'_{i,r} \geq a_{i,r}$, since $a_{i,r} = 0$ for these rounds by the definition of r_i . Finally,

$$\begin{aligned} U_{i,R}(\mathbf{a}_i) - U_{i,R}(\mathbf{a}'_i) &= \sum_{r=1}^R (u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r})) \\ &= (u_{i,r'}(a_{i,r'}) - u_{i,r'}(a'_{i,r'})) + \sum_{r \neq r'} (u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r})) \\ &= H(a_{i,r'} - a'_{i,r'}) - \sum_{r \neq r'} (u_{i,r}(a'_{i,r}) - u_{i,r}(a_{i,r})) \\ &\geq H(a_{i,r'} - a'_{i,r'}) - H(a_{i,r'} - a'_{i,r'}) = 0 \end{aligned}$$

Here, the third transition follows from Lemma 18, and the final transition follows because $\sum_{r \neq r'} (a'_{i,r'} - a_{i,r}) = a_{i,r'} - a'_{i,r'}$, and every term in the sum is positive.

The proof for the case where $d'_{i,r'} > d_{i,r'}$ is in the Appendix. Together, they show that i achieves (weakly) higher utility by truthfully reporting her demand $d_{i,r'}$ at round r , rather than misreporting $d'_{i,r'} \neq d_{i,r'}$. By the argument at the start of this subsection, this is sufficient to prove strategy-proofness. \square

5.4 Approximating Sharing Incentives

Unfortunately, FL fails to satisfy SI, and may give an agent as little as half of her SI share.

THEOREM 22. *FL does not satisfy α -SI for any $\alpha > 0.5$.*

PROOF. Consider an instance with R rounds, and $R + 1$ agents, each with endowment $e_i = 1$. Agent 1 has $d_{1,1} = d_{1,R} = 1$ and $d_{1,2} = \dots = d_{1,R-1} = 0$, agent 2 has $d_{2,r} = R$ for all rounds r , and all other agents have $d_{i,r} = 0$ for all rounds r . In round 1, agent 1 receives allocation $a_{1,1} = 1$ and agent 2 receives $a_{2,1} = R$. For rounds $r = 2, \dots, R - 1$, each agent $j \neq 2$ receives allocation $a_{j,r} = 1 + 1/R$. Therefore, in round R , agent 1 receives $a_{1,R} = R - 1 - (R - 2)(1 + 1/R) = 2/R$. Her total utility is therefore $((R + 2)/R)H + (R - (R + 2)/R)L$, compared to total utility $2H + (R - 2)L$ that she would have received by not participating in the mechanism. For $L = 0$, the ratio of these utilities approaches 0.5 as $R \rightarrow \infty$. \square

However, FL does provide a 0.5 approximation guarantee to SI, as we show in the remainder of this subsection. We suppose that agent i truthfully reports her demand $d_{i,r}$ for all rounds (since FL is SP, she could do no better by lying), and show that she receives at least half as much utility as she would by not participating.

Recall that for every agent i , we denote by r_i the first round at which $a_{i,r_i} = t_{i,r_i} > 0$. For every agent i , define sets B_i and A_i to be the agents that run out of tokens before and after i , respectively. Formally,

$$B_i = \{j : r_j \leq r_i \text{ and } a_{j,r_i}/e_j < a_{i,r_i}/e_i\}$$

$$A_i = \{j : r_j \geq r_i \text{ and } r_j = r_i \implies a_{j,r_i}/e_j \geq a_{i,r_i}/e_i\}.$$

For a round r , define

$$s_{i,r} = a_{i,r} - e_i \frac{\sum_{j \in A_i} a_{j,r}}{\sum_{j \in A_i} e_j}.$$

That is, $s_{i,r}$ is the number of resources i gets more than the (endowment weighted) average number of resources for agents in A_i . Note further that

$$\sum_{r=1}^R s_{i,r} = \sum_{r=1}^R a_{i,r} - \frac{e_i}{\sum_{j \in A_i} e_j} \sum_{j \in A_i} \sum_{r=1}^R a_{j,r} = e_i - \frac{e_i}{\sum_{j \in A_i} e_j} \sum_{j \in A_i} e_j = 0.$$

LEMMA 23. *For every agent i and every round r , $s_{i,r} \leq \min(d_{i,r}, a_{i,r})$.*

PROOF. If $a_{i,r} \leq d_{i,r}$, then the lemma statement says that $s_{i,r} \leq a_{i,r}$, which is obviously true by the definition of $s_{i,r}$. If $a_{i,r} > d_{i,r}$, then we know from the definition of FL that $\sum_{j \in [n]} \min(d_{j,r}, t_{j,r}) < E$, and $a_{i,r} = \min(xe_i, t_{i,r})$, where x is the objective value of FL's call to the PSWC program. Further, all agents with $\frac{a_{j,r}}{e_j} < \frac{a_{i,r}}{e_i} \leq x$ are those with $a_{j,r} = t_{j,r}$, so by definition, $r_j \leq r_i$ and $\frac{a_{j,r}}{e_j} < \frac{a_{i,r}}{e_i}$, which means $j \in B_i$. Therefore, $\frac{a_{j,r}}{e_j} \geq \frac{a_{i,r}}{e_i}$ for all $j \in A_i$, which implies $\frac{\sum_{j \in A_i} a_{j,r}}{\sum_{j \in A_i} e_j} \geq \frac{a_{i,r}}{e_i}$. To complete the proof, note that

$$s_{i,r} = a_{i,r} - e_i \frac{\sum_{j \in A_i} a_{j,r}}{\sum_{j \in A_i} e_j} \leq a_{i,r} - e_i \frac{a_{i,r}}{e_i} = 0 \leq d_{i,r} = \min(d_{i,r}, a_{i,r}).$$

□

THEOREM 24. *Under FL, agents receive at least half the number of high-valued resources that they would have received under static allocations.*

PROOF. Let S denote the number of high-valued resources that agent i receives under static allocations. While i has tokens remaining, under FL, she is guaranteed to get as many resources as she demands up to her endowment e_i . Thus, for these rounds, she would obtain no additional high-valued resources from not participating in the mechanism. However, there is the possibility that by participating in the mechanism, she runs out of tokens prematurely, thus missing out on resources in later rounds that she wants, and would have received by not participating in the mechanism (as in the proof of Theorem 22). The proof proceeds by showing that for every resource that i does not receive due to a lack of tokens, she must have received at least one high-valued resource in an earlier round.

Suppose first that $a_{i,r_i} \geq e_i$. We have the following inequality:

$$\sum_{r \leq r_i} \min(d_{i,r}, a_{i,r}) \geq \sum_{r \leq r_i} s_{i,r} = - \sum_{r > r_i} s_{i,r} = \sum_{r > r_i} \left(e_i \frac{\sum_{j \in A_i} a_{j,r}}{\sum_{j \in A_i} e_j} \right) = \sum_{r > r_i} \left(\frac{E}{\sum_{j \in A_i} e_j} \right) e_i \geq (T - r_i) e_i. \quad (1)$$

The first inequality follows from Lemma 23, and the second inequality because $\sum_{j \in A_i} e_j \leq E$. The first equality holds because $\sum_{r=1}^R s_{i,r} = 0$, and the second equality holds because $a_{i,r} = 0$ for all $r > r_i$. The third equality holds because for rounds $r > r_i$, only agents in A_i remain active, so all resources are allocated to them.

Note that S , the number of high-valued resources that i receives by not sharing, is upper bounded by

$$\begin{aligned} S &\leq \sum_{r=1}^R \min(d_{i,r}, e_i) \leq \sum_{r \leq r_i} \min(d_{i,r}, e_i) + \sum_{r > r_i} e_i \\ &\leq \sum_{r \leq r_i} \min(d_{i,r}, a_{i,r}) + \sum_{r > r_i} e_i \\ &= \sum_{r \leq r_i} \min(d_{i,r}, a_{i,r}) + (T - r_i) e_i \\ &\leq 2 \sum_{r \leq r_i} \min(d_{i,r}, a_{i,r}). \end{aligned}$$

The third inequality holds because under FL guarantees each agent $\min(d_{i,r}, e_i)$ resources, provided they have sufficient tokens remaining, which is the case because we assume $a_{i,r_i} \geq e_i$. The final inequality follows from Equation (1). Since agent i receives exactly $\sum_{r \leq r_i} \min(d_{i,r}, a_{i,r}) \geq S/2$ resources from participating in FL, the lemma holds in this case.

Second, suppose that $a_{i,r_i} < e_i$. We have the following inequality:

$$\begin{aligned} \sum_{r \leq r_i} \min(d_{i,r}, a_{i,r}) &\geq \sum_{r < r_i} \min(d_{i,r}, a_{i,r}) \geq \sum_{r < r_i} s_{i,r} = - \sum_{r > r_i} s_{i,r} - s_{i,r_i} \\ &\geq e_i(T - r_i) + e_i \frac{\sum_{j \in A_i} a_{j,r_i}}{\sum_{j \in A_j} e_j} - a_{i,r_i} \geq e_i(T - r_i) + e_i - a_{i,r_i} = e_i(T - r_i + 1) - a_{i,r_i} \quad (2) \end{aligned}$$

The first inequality holds because $\min(d_{i,r_i}, a_{i,r_i}) \geq 0$. The second inequality follows from Lemma 23, and the third inequality holds from Equation (1) and the definition of s_{i,r_i} . The fourth inequality holds because at round r_i , agent i receives allocation $a_{i,r_i} < e_i$, therefore every agent $j \in B_i$ receives allocation $a_{j,r_i} < e_j$, therefore $\sum_{j \in A_i} a_{j,r_i} \geq \sum_{j \in A_i} e_j$.

As with the previous case, we can derive an upper bound on S , the number of high-valued resources i would receive by not sharing. First, suppose that $a_{i,r_i} > d_{i,r_i}$. Then we have

$$\begin{aligned}
S &\leq \sum_{r=1}^R \min(d_{i,r}, e_i) \leq \sum_{r < r_i} \min(d_{i,r}, e_i) + d_{i,r_i} + \sum_{r > r_i} e_i \\
&\leq \sum_{r < r_i} \min(d_{i,r}, a_{i,r}) + \min(d_{i,r_i}, a_{i,r_i}) + \sum_{r > r_i} e_i \\
&= \sum_{r \leq r_i} \min(d_{i,r}, a_{i,r}) + (T - r_i)e_i \\
&\leq \sum_{r \leq r_i} \min(d_{i,r}, a_{i,r}) + (T - r_i + 1)e_i - a_{i,r_i} \\
&\leq 2 \sum_{r \leq r_i} \min(d_{i,r}, a_{i,r})
\end{aligned}$$

The third inequality holds because FL guarantees each agent $\min(d_{i,r}, e_i)$ resources, provided they have sufficient tokens remaining, and by the assumption that $a_{i,r_i} > d_{i,r_i}$, the fourth inequality from the assumption that $a_{i,r_i} < e_i$, and the final inequality from Equation (2). Next, suppose that $a_{i,r_i} \leq d_{i,r_i}$. Then we have

$$\begin{aligned}
S &\leq \sum_{r=1}^R \min(d_{i,r}, e_i) \leq \sum_{r < r_i} \min(d_{i,r}, e_i) + e_i + \sum_{r > r_i} e_i \\
&\leq \sum_{r < r_i} \min(d_{i,r}, a_{i,r}) + a_{i,r_i} + (e_i - a_{i,r_i}) + \sum_{r > r_i} e_i \\
&= \sum_{r \leq r_i} \min(d_{i,r}, a_{i,r}) + (T - r_i + 1)e_i - a_{i,r_i} \\
&\leq 2 \sum_{r \leq r_i} \min(d_{i,r}, a_{i,r})
\end{aligned}$$

The third inequality holds because FL guarantees each agent $\min(d_{i,r}, e_i)$ resources, provided they have sufficient tokens remaining, the equality from the assumption that $a_{i,r_i} \leq d_{i,r_i}$, and the final inequality from Equation (2).

As with the previous case, $e_i(T - r_i + 1) - a_{i,r_i}$ is an upper bound on the number of H valued resources that i may have been able to receive in rounds $r \geq r_i$ had she not participated in the mechanism, over and above those she receives by participating. $\sum_{r \leq r_i} \min(d_{i,r}, a_{i,r})$ is the number of H valued resources she receives by participating in the mechanism. Therefore $\sum_{r \leq r_i} \min(d_{i,r}, a_{i,r}) + e_i(T - r_i + 1) - a_{i,r_i} \leq 2 \sum_{r \leq r_i} \min(d_{i,r}, a_{i,r})$ is an upper bound on the number of H valued resources i would receive by not participating in the mechanism. Therefore, i receives at least half as many H valued resources from participating as she would have by not participating. \square

Note that Theorem 24 implies the desired approximation. Suppose that i obtains utility $SH + (Re_i - S)L$ by not participating in the mechanism. Theorem 24 in combination with the fact that she will receive the same number of resources overall whether she participates or not, implies that, by participating, she gets at least $SH/2 + (Re_i - S/2)L \geq SH/2 + (Re_i/2 - S/2)L = (SH + (Re_i - S)L)/2$.

5.5 Limit Efficiency for Symmetric Agents

In this section, we prove that, under certain assumptions, FL is efficient in the limit as the number of rounds grows large. Suppose that each agent has the same endowment. Without loss of generality,

suppose that each agent has $e_i = 1$. Further, suppose that demands are drawn i.i.d. across rounds and that the distribution within rounds treats agents symmetrically, either demands are drawn i.i.d. across agents, or there is correlation that treats all agents symmetrically.

THEOREM 25. *When demands are drawn i.i.d. across rounds and agents are symmetric, FL achieves an $(R - R^{2/3})/R$ fraction of the optimal efficiency with probability at least $1 - n^3/R^{1/3}$. In particular, FL approaches full efficiency with high probability in the limit as the number of rounds grows large.*

PROOF. Suppose we are in a world where tokens are unlimited. Let Q be a random variable denoting how many tokens a single agent i would spend (i.e. how many resources i would be allocated) in a single round. Note that Q can never take a value larger than n , since only n resources are allocated per round. Note that by the symmetry of the agents, Q is independent of the identity of any single agent, and independent of the particular round since FL allocates independently of the round. By symmetry, $\mathbb{E}(Q) = 1$. Let $\text{StdDev}(Q) = \sigma \leq n$, where the inequality holds because Q is bounded by n . Let $r = R - R^{2/3}$ and let Q_r be a random variable denoting the number of tokens i would spend before the start of round $r + 1$. Because demands are drawn independently across rounds, and no agent runs out of tokens, $\mathbb{E}(Q_r) = r$ and $\text{StdDev}(Q_r) = \sqrt{r}\sigma$.

Consider the probability that agent i spends at least R tokens in the first r rounds:

$$\begin{aligned} P(Q_r \geq R) &= P(Q_r - \mathbb{E}(Q_r) \geq R - r) \\ &= P(Q_r - \mathbb{E}(Q_r) \geq R^{2/3}) \\ &= P(Q_r - \mathbb{E}(Q_r) \geq \frac{R^{1/6}}{\sigma} \sqrt{R}\sigma) \\ &\leq P(Q_r - \mathbb{E}(Q_r) \geq \frac{R^{1/6}}{\sigma} \sqrt{r}\sigma) \\ &\leq \frac{\sigma^2}{R^{1/3}} \end{aligned}$$

Here the final inequality follows from Chebyshev's concentration inequality, because $\sqrt{r}\sigma$ is the standard deviation of Q_r . Taking a union bound over all n agents, the probability that any agent spends at least R tokens in the first r rounds is at most $n\sigma^2/R^{1/3} \leq n^3/R^{1/3}$.

Now suppose agents are limited by R tokens. If some agent runs out of tokens in this world, then it must also be the case that some agent spent at least R tokens in the unlimited token world. Therefore, the probability that any agent runs out of tokens is at most the probability that some agent spends more than R tokens in the unlimited token world, which is at most $n^3/R^{1/3}$. This approaches 0 as $R \rightarrow \infty$. So, with probability going to 1, no agent runs out of tokens before round r .

By the definition of FL, full efficiency is achieved on all rounds for which no agents have their allocation limited by lack of tokens. Therefore, with probability going to 1, FL allocates efficiently for the first r rounds. Therefore, because demands are i.i.d. across rounds, the expected efficiency of the mechanism approaches at least an $r/R = (R - R^{2/3})/R$ fraction of the optimal efficiency. This fraction approaches 1 as $R \rightarrow \infty$. \square

6 T-PERIOD MECHANISM

We have shown that FL satisfies strategy-proofness and a theoretical asymptotic efficiency guarantee. Further, as we show in §7, FL exhibits only small efficiency loss in practice in settings where our theoretical guarantee does not apply. However, FL does not achieve (full) sharing incentives. In settings where agents require a strong guarantee in order to participate, it may be desirable to strictly enforce sharing incentives, in which case FL is not a suitable choice. In this section, we introduce the *T-Period mechanism*, which satisfies both SP and SI. While the *T-Period mechanism*

does exhibit some gains from sharing (i.e., is more efficient than static allocation), it sacrifices some efficiency relative to FL.

6.1 Definition

The T -Period mechanism splits the rounds into periods of length $2T$.¹ For the first T rounds of each period, we allow the agents to ‘borrow’ unwanted resources from others. In the last T rounds of each period, the agents ‘pay back’ the resources so that their cumulative allocation across the entire period is equal to their endowment, $2Te_i$.

The allocations in the second set of T rounds are independent of reports and determined completely by the allocations in the first set of T rounds. Note that because the number of resources that an agent i can pay back over T rounds is bounded by Te_i , we allow an agent to borrow at most Te_i resources (i.e., receive at most $2Te_i$ resources) over the first T rounds of a period.

Algorithm 2 T -Period Mechanism

Input. Agents’ reported demands, \mathbf{d}' , and their endowments, \mathbf{e}

Output. Agents’ allocations, \mathbf{a}

```

for  $r \in \{1, \dots, R\}$  do
  if  $(r \bmod 2T) = 1$  then
     $\mathbf{b} \leftarrow T\mathbf{e}$                                  $\triangleright b_i$  is the amount that  $i$  is able to borrow
     $\mathbf{y} \leftarrow \mathbf{0}$                                  $\triangleright$  resources received so far this period.
  end if
  if  $1 \leq (r \bmod 2T) \leq T$  then
     $\bar{\mathbf{d}} \leftarrow \min(\mathbf{d}'_{\cdot,r}, \mathbf{e} + \mathbf{b})$             $\triangleright \bar{d}_i$  is  $i$ 's allocatable demand
     $D \leftarrow \sum_{i \in [n]} \bar{d}_i$ 
    if  $D \geq E$  then
       $\mathbf{a}_{\cdot,r} \leftarrow \text{PSWC}(A = E, \mathbf{l} = \bar{\mathbf{d}}, \mathbf{m} = \mathbf{0}, \mathbf{w} = \mathbf{e})$ 
    else
       $\mathbf{a}_{\cdot,r} \leftarrow \text{PSWC}(A = E, \mathbf{l} = \mathbf{e} + \mathbf{b}, \mathbf{m} = \bar{\mathbf{d}}, \mathbf{w} = \mathbf{e})$ 
    end if
     $\mathbf{y} \leftarrow \mathbf{y} + \mathbf{a}_{\cdot,r}$ 
     $\mathbf{b} \leftarrow \mathbf{b} - \max(\mathbf{0}, \mathbf{a}_{\cdot,r} - \mathbf{e})$ 
  else
     $\mathbf{a}_{\cdot,r} \leftarrow \frac{1}{T}(2T\mathbf{e} - \mathbf{y})$ 
  end if
end for

```

In Algorithm 2, each agent i has a borrowing limit, b_i , which is defined to be the maximum amount of resources that agent i can borrow in whatever remains of the first T rounds of each period. For our analysis, we denote the value of b_i at the start of round r by $b_{i,r}$. At the beginning of each period, we set $b_{i,r}$ to be Te_i , because agent i can at most pay back her whole endowment, e_i , at every T ‘payback’ rounds. We again define \bar{d}_i to be the allocatable demand of agent i at each round of the first T rounds and refer to $\bar{d}'_{i,r}$ as agent i 's allocatable demand at round r . At each round r , the allocatable demand of agent i is the minimum of her reported demand $d'_{i,r}$, and her endowment plus her borrowing limit, $e_i + b_{i,r}$.

¹For convenience, we suppose that R is a multiple of $2T$. If this is not the case, we can adapt the mechanism by returning each agent their endowment for any leftover rounds.

We illustrate the T -Period mechanism with an example.

EXAMPLE 26. Consider the instance from Example 12, where each agent has endowment $e_i = 1$ and demands are given by:

	$d_{i,1}$	$d_{i,2}$	$d_{i,3}$	$d_{i,4}$
$i = 1$	3	1	1	0
$i = 2$	0	2	1	2
$i = 3$	0	0	0	4

When $T = 1$, agents can ‘borrow’ resources at odd rounds and ‘pay back’ those resources at even rounds. Therefore, the maximum allocatable demand for each agent and at each round is 2, because the ‘payback’ period only has one round. The 1-Period (1-P) mechanism allocates resources as follows.

	$a_{i,1}^{1-P}$	$a_{i,2}^{1-P}$	$a_{i,3}^{1-P}$	$a_{i,4}^{1-P}$
$i = 1$	2	0	1	1
$i = 2$	0.5	1.5	1	1
$i = 3$	0.5	1.5	1	1

At round 1, agent 1 wants 2 extra resources in addition to her endowment. However, under 1-P, she can only afford 1 extra resource. She borrows 0.5 resources from agent 2 and 0.5 resources from agent 3. At round 2, agent 1 pays back agents 2 and 3 and receives zero resources. When $T = 1$, the mechanism rigidly forces agents to pay back resources right after they borrow them. Agent 1 would prefer to get her high-valued resource at round 2 and delay paying back agents 2 and 3 until the last round where her demand is zero. Note that agent 3 would also prefer to be paid back in the last round, the only round in which she has non-zero demand.

To see how increasing T allows more flexibility, consider $T = 2$ for the same example. The 2-Period (2-P) mechanism allocates resources as follows.

	$a_{i,1}^{2-P}$	$a_{i,2}^{2-P}$	$a_{i,3}^{2-P}$	$a_{i,4}^{2-P}$
$i = 1$	3	1	0	0
$i = 2$	0	2	1	1
$i = 3$	0	0	2	2

Agent 2 is allowed to borrow 2 extra resources over the first two rounds, whereas, under the 1-P mechanism, she is never allowed to borrow more than one resource per round. She borrows these two resources at the first round from agents 2 and 3, and pays them back at rounds 3 and 4.

Since the T -Period mechanism increases flexibility over the static mechanism, it provides some gains from sharing. We would expect that increasing T , in general, will improve efficiency as it allows for ‘borrowed’ resources to be spent more flexibly. In the following subsection, we show that these efficiency gains do not harm SI or SP when $T \leq 2$. Many proofs closely follow those in §5 and are deferred to the Appendix.

6.2 Axiomatic Properties of T -Period Mechanism

We first state a lemma characterizing the allocations of the T -Period mechanism that is analogous to Lemma 13

LEMMA 27. Let x denote the objective value of a call to PSWC. Suppose that $1 \leq (r \bmod 2T) \leq T$. If $D \geq E$, then $a_{i,r} = \min(e_i + b_i, d'_{i,r}, xe_i)$. If $D < E$, then $a_{i,r} = \min(e_i + b_i, \max(d'_{i,r}, xe_i))$.

To prove strategy-proofness of the 1-Period and 2-Period mechanisms, we show that no agent has an incentive to report $d'_{i,r} \neq d_{i,r}$ for any round r . We again consider parallel cases, one in which

agent i misreports $d'_{i,r}$ and one in which she truthfully reports $d_{i,r}$ with all other reports the same across the two cases. Allocations and borrowing limits in the former case is denoted by a' and b' respectively, and by a and b in the latter case. Let D_r denote the total allocatable demand at a round r in the truthful case and D'_r denote the total allocatable demand at a round r in the misreported case.

Since the T -Period mechanism resets every $2T$ rounds, we can assume without loss of generality that $R = 2T$ for the sake of reasoning about SP and SI. For rounds $r > T$, the allocations depend completely on the allocations at earlier rounds, and not on the agents' reports, so there is clearly no benefit to an agent for misreporting in these rounds. It remains to show that reporting $d'_{i,r} = d_{i,r}$ is optimal for rounds $r \leq T$.

Our next lemma is analogous to Lemma 16.

LEMMA 28. *Let $a_{i,r}$ and $a'_{i,r}$ denote the allocations of agent i at round r when she reports $d_{i,r}$ and $d'_{i,r}$, respectively, holding fixed the reports of all agents $j \neq i$ and agent i 's reports on all rounds other than r . If $d'_{i,r} < d_{i,r}$ then $a'_{i,r} \leq a_{i,r}$, and $a'_{j,r} \geq a_{j,r}$ for all $j \neq i$.*

Suppose that i reports $d'_{i,r} \neq d_{i,r}$ for some round r , but this misreport does not change i 's allocation (that is, $a'_{i,r} = a_{i,r}$). By Lemma 28, $a'_{j,r} = a_{j,r}$ for all $j \neq i$. Therefore, i 's misreport has not changed the allocations at round r . Since all future rounds take into account allocations at previous rounds but not reports, i 's misreport has had no effect on the allocations in any round. Thus, i did not benefit from this misreport. We therefore assume that $a'_{i,r} \neq a_{i,r}$ for any round r where i reports $d'_{i,r} \neq d_{i,r}$ in the remainder of this section.

The next lemma and corollary are analogous to Lemma 17 and Corollary 18. They say that if i obtains fewer resources from misreporting at round r , then those resources are all high-valued resources.

LEMMA 29. *Hold the reports of all agents $j \neq i$ fixed, and the reports of agent i on all rounds other than r fixed. If i reports $d'_{i,r} < d_{i,r}$ and receives $a'_{i,r} < a_{i,r}$, then $a_{i,r} \leq d_{i,r}$.*

As a corollary we obtain a formula for the difference between the utility that agent i receives at round r under truthful reporting and misreporting, when i gets fewer resources in the misreported instance.

COROLLARY 30. *Hold the reports of all agents $j \neq i$ fixed, and the reports of agent i on all rounds other than r fixed. If i reports $d'_{i,r} < d_{i,r}$ and receives $a'_{i,r} < a_{i,r}$, then $u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r}) = H(a_{i,r} - a'_{i,r})$.*

The next lemma and corollary complement Lemma 29 and Corollary 30 in the case where i receives more resources in the misreported instance than the truthful instance at round r .

LEMMA 31. *Hold the reports of all agents $j \neq i$ fixed, and the reports of agent i on all rounds other than r fixed. If i reports $d'_{i,r} > d_{i,r}$ and receives $a'_{i,r} > a_{i,r}$, then $a_{i,r} \geq d_{i,r}$.*

COROLLARY 32. *Hold the reports of all agents $j \neq i$ fixed, and the reports of agent i on all rounds other than r fixed. If i reports $d'_{i,r} > d_{i,r}$ and receives $a'_{i,r} > a_{i,r}$, then $u_{i,r}(a'_{i,r}) - u_{i,r}(a_{i,r}) = L(a'_{i,r} - a_{i,r})$.*

We can now show that misreporting in round T is never beneficial to an agent.

LEMMA 33. *An agent never improves her utility by reporting $d'_{i,T} \neq d_{i,T}$.*

As a corollary, we immediately have that the 1-Period mechanism is strategy-proof, because misreporting at round $r = 1 = T$ is not beneficial, and misreporting at round $r = 2 > T$ is not beneficial by our earlier argument.

COROLLARY 34. *The 1-Period mechanism satisfies strategy-proofness.*

Our next lemma is a monotonicity statement for the borrowing limits: if i 's borrowing limit at round r increases, and all other agents' borrowing limits decrease, then i 's allocation (weakly) increases and all other agents' allocations (weakly) decrease.

LEMMA 35. *Suppose that $r \leq T$. If $b'_{i,r} \geq b_{i,r}$ and $b'_{j,r} \leq b_{j,r}$ for all $j \neq i$, and $d'_{k,r} = d_{k,r}$ for all agents k , then $a'_{i,r} \geq a_{i,r}$.*

We now show that the 2-Period mechanism is strategy-proof.

THEOREM 36. *The 2-Period mechanism satisfies strategy-proofness.*

PROOF. By Lemma 33, no agent can benefit by reporting $d'_{i,2} \neq d_{i,2}$. Similarly, no agent can benefit by reporting $d'_{i,r} \neq d_{i,r}$ for $r \in \{3, 4\}$, because the 2-Period mechanism ignores reports for those rounds. We may therefore assume that $d'_{i,r} = d_{i,r}$ for all agents i and all rounds $r \geq 2$.

We show that an agent cannot benefit from reporting $d'_{i,1} < d_{i,1}$. The proof that reporting $d'_{i,1} > d_{i,1}$ is not beneficial is very similar. If $a'_{i,1} = a_{i,1}$, then $a'_{j,1} = a_{j,1}$ for all $j \neq i$, by Lemma 28. Therefore, the allocations are unchanged for all rounds i , as the 2-Period mechanism takes into account allocations at earlier rounds, but not reports, and the allocations at round 1 are the same in the truthful and misreported instances. We therefore assume that $a_{i,1} = a'_{i,1} + k$, for some $k > 0$. This implies that $b_{i,2} = b'_{i,2} - k_i$, for some $k_i \leq k$. By Corollary 30, i receives kH more utility in round 1 under truthful reporting than under misreporting. For every $j \neq i$, $a_{j,1} \leq a'_{j,1}$, and $b_{j,2} = b'_{j,2} + k_j$, where $\sum_{j \neq i} k_j \leq k$. By Lemma 35, $a'_{i,2} \geq a_{i,2}$. In the following, we show that $a'_{i,2} \leq a_{i,2} + k$. Let x and x' denote the objective value in the T-Period mechanism's call to PSWC when i reports $d_{i,r}$ and $d'_{i,r}$, respectively. We consider four cases, corresponding to whether resources in the truthful and misreported instances are over or under demanded at round 2. Suppose first that $D_2 \geq E$ and $D'_2 \geq E$. First, suppose that $x' < x$. Then, by Lemma 27,

$$\begin{aligned} a'_{i,2} &= \min(e_i + b'_i, d_{i,2}, x'e_i) = \min(e_i + b_{i,2} + k_i, d_{i,2}, x'e_i) \\ &\leq \min(e_i + b_{i,2}, d_{i,2}, x'e_i) + k_i \leq \min(e_i + b_{i,2}, d_{i,2}, xe_i) + k_i \leq a_{i,2} + k \end{aligned}$$

Next, suppose that $x' \geq x$. Then for all $j \neq i$,

$$\begin{aligned} a'_{j,2} &= \min(e_j + b'_j, d_{j,2}, x'e_j) = \min(e_j + b_{j,2} - k_j, d_{j,2}, x'e_j) \\ &\geq \min(e_j + b_{j,2}, d_{j,2}, x'e_j) - k_j \geq \min(e_j + b_{j,2}, d_{j,2}, xe_j) - k_j = a_{j,2} - k_j \end{aligned}$$

Taking the sum over all $j \neq i$ and noting that $\sum_{j \neq i} k_j \leq k$, we have that $\sum_{j \neq i} a'_{j,2} \geq \sum_{j \neq i} a_{j,2} - k$. Therefore, $a'_{i,2} \leq a_{i,2} + k$. Second, suppose that $D_2 \geq E$ and $D'_2 < E$. Then, by the definition of the T-Period mechanism, $a_{j,2} \leq \min(e_j + b_{j,2}, d_{j,2})$ for all $j \neq i$. Further

$$a'_{j,2} \geq \min(e_j + b'_j, d_{j,2}) = \min(e_j + b_{j,2} - k_j, d_{j,2}) \geq \min(e_j + b_{j,2}, d_{j,2}) - k_j \geq a_{j,2} - k_j$$

By the same argument as in the previous case, this implies that $a'_{i,2} \leq a_{i,2} + k$. Third, suppose that $D_2 < E$ and $D'_2 \geq E$. Then

$$a'_{i,2} \leq \min(e_i + b'_i, d_{i,2}) = \min(e_i + b_{i,2} + k_i, d_{i,2}) \leq \min(e_i + b_{i,2}, d_{i,2}) + k_i \leq a_{i,r} + k$$

Finally, suppose that $D_2 < E$ and $D'_2 < E$. First, suppose that $x' < x$. Then

$$\begin{aligned} a'_{i,2} &= \min(e_i + b'_i, \max(d_{i,2}, x'e_i)) = \min(e_i + b_{i,2} + k_i, \max(d_{i,2}, x'e_i)) \\ &\leq \min(e_i + b_{i,2}, \max(d_{i,2}, x'e_i)) + k_i \\ &\leq \min(e_i + b_{i,2}, \max(d_{i,2}, xe_i)) + k_i \leq a_{i,2} + k \end{aligned}$$

Next, suppose that $x' \geq x$. Then for all $j \neq i$,

$$\begin{aligned} a'_{j,2} &= \min(e_j + b'_j, \max(d_{j,2}, x'e_j)) = \min(e_j + b_{j,2} - k_j, \max(d_{j,2}, x'e_j)) \\ &\geq \min(e_j + b_{j,2}, \max(d_{j,2}, x'e_j)) - k_j \\ &\geq \min(e_j + b_{j,2}, \max(d_{j,2}, xe_j)) - k_j = a_{j,2} - k_j \end{aligned}$$

Again, this implies that $a'_{i,2} \leq a_{i,2} + k$.

In all cases, we have that $a_{i,2} \leq a'_{i,2} \leq a_{i,2} + k$. Therefore, $a'_{i,1} + a'_{i,2} \leq a_{i,1} + a_{i,2}$, which means that $a'_{i,3} \geq a_{i,3}$ and $a'_{i,4} \geq a_{i,4}$. Consider the difference in utility across all four rounds between the truthful and misreported instances.

$$\begin{aligned} U_{i,4}(\mathbf{a}_i) - U_{i,4}(\mathbf{a}'_i) &= \sum_{r=1}^4 (u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r})) \\ &= kH + \sum_{r=2}^4 (u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r})) \geq kH - kH = 0 \end{aligned}$$

The second transition is by Corollary 30, and the third transition because each $a'_{i,r} \geq a_{i,r}$ for all $r \in \{2, 3, 4\}$, $\sum_{r=2}^4 (a'_{i,r} - a_{i,r}) = k$, and each resource can be worth at most H to agent i . \square

Given that the 1-P and 2-P mechanisms satisfy SP, it is easy to see that they satisfy SI also. By strategy-proofness, the utility that an agent gets from truthfully reporting her demands is at least the utility she gets from reporting $d'_{i,r} = e_i$ for all rounds r . Sharing incentives therefore follows as a corollary of the following proposition.

PROPOSITION 37. *Under the T-Period mechanism, any agent that reports $d'_{i,r} = e_i$ for all rounds r receives $a_{i,r} = e_i$ for all rounds r .*

COROLLARY 38. *The T-Period mechanism satisfies SI for $T \leq 2$.*

One may hope to continue increasing flexibility, and therefore performance, by increasing the length of the ‘borrowing’ and ‘payback’ periods, potentially all the way to having a single borrowing period of length $R/2$ and a single payback period of length $R/2$. Unfortunately, even for periods of length 3, strategy-proofness is violated.

EXAMPLE 39. *Consider the 3-P mechanism. Suppose that $n = 5$ and $R = 6$. Each agent has endowment $e_i = 1$ (so each agent can borrow a total of three resources over the first three rounds, corresponding to the sum of their endowment across the final three rounds). Truthful demands are given by the following table.*

	$d_{i,1}$	$d_{i,2}$	$d_{i,3}$	$d_{i,4}$	$d_{i,5}$	$d_{i,6}$
$i = 1$	3	3	0	1	1	1
$i = 2$	0	3	3	1	1	1
$i = 3$	0	0	0	0	0	0
$i = 4$	0	0	0	0	0	0
$i = 5$	0	0	0	0	0	0

The corresponding allocations are given by:

	$a_{i,1}^{3-P}$	$a_{i,2}^{3-P}$	$a_{i,3}^{3-P}$	$a_{i,4}^{3-P}$	$a_{i,5}^{3-P}$	$a_{i,6}^{3-P}$
$i = 1$	3	2	0.75	0.08	0.08	0.08
$i = 2$	0.5	3	2	0.17	0.17	0.17
$i = 3$	0.5	0	0.75	1.58	1.58	1.58
$i = 4$	0.5	0	0.75	1.58	1.58	1.58
$i = 5$	0.5	0	0.75	1.58	1.58	1.58

Agent 1's utility is $5.25H + 0.75L$. If agent 1 misreports $d'_{1,1} = 2$, it can be checked that her allocations become 2, 2.5, 0.625, 0.292, 0.292, 0.292. Her utility is then $5.375H + 0.625L$, which is higher than her utility from reporting truthfully.

7 EVALUATION

In this section, we evaluate different mechanisms using real and synthetic benchmarks. For real benchmarks, we use a Google cluster trace [1, 27], which data collected from a 12.5k-machine cluster over a month-long period in May 2011. All the machines in the cluster share a common cluster manager that allocates agent tasks to machines.

Agents submit a set of resource demands for each task (e.g. required processors, memory, or disk space). Agent demands are normalized relative to the largest capacity of the resource on any machine in the traces. The cluster manager records any changes in the status of tasks (e.g. being evicted, failed, or killed) during their life cycle in a task event table. We use the task event table to track agents' demands for processors over time. Note that since all demands are scaled by the same factor, we safely use normalized demands as actual demands.

We divide time into 15 min intervals.² We define agents' demands for each interval to be the sum of their demands for all tasks they run in that interval. After processing the traces, we remove agents with constant demands or with average demand less than some marginal threshold. We assume that agents' endowments are equal to their average demands.

We observe that, for each agent, demands computed from Google traces have high correlations over time. An agent with high demand at 12am has typically high demand at 12:15am as well. In some deployment scenarios, demands may not be highly correlated. For example, when university cluster machines are allocated to professors and researchers on a daily basis, a researcher may have some jobs today, but may not want to use the cluster tomorrow.

To evaluate mechanisms in scenarios without correlated demands, we use synthetic benchmarks. We create random agent populations and random number of rounds. For each agent, we uniformly and randomly assign an endowment from 1 to 20. Once agents' endowments are set, we uniformly and randomly generate agent demands such that their average is equal to agents' endowments (i.e. $d_{i,r} \sim u[0, 2e_i]$)

Metrics. We report social welfare and Nash welfare, focusing on the number of high-valued resources that each mechanism allocates. For social welfare, we report the following.

$$\text{Social Welfare} = \sum_i \sum_r \min(d_{i,r}, a_{i,r}).$$

Social welfare is a measure of efficiency but fails to distinguish between fair and unfair outcomes. For instance, suppose agent A with endowment 100 and agent B with endowment 1 both have demand 101. Allocating 100 units to agent A and 1 unit to agent B has the same social welfare as allocating 1 unit to A and 100 units to B. To distinguish between these two allocations, we also report the (weighted) Nash welfare as follows.

²We have created demands for varying time intervals. Since results do not change significantly for different interval lengths, we only include results on 15-min-long intervals.




Fig. 3. Normalized Social Welfare. Social welfare achieved by different dynamic allocation mechanisms normalized to that of static allocations for Google cluster traces and 100 instances of random demands.




Fig. 4. Normalized Nash Welfare. The Nash welfare achieved by different dynamic allocation mechanisms normalized to that of static allocations for Google cluster traces and 100 instances of random demands.




Fig. 5. Sensitivity of the Flexible Lending Mechanism. Social welfare of the flexible lending mechanism normalized to that of static allocations for varying agent population sizes and numbers of rounds. We fix the number of rounds to 50 when we vary the number of agents, and fix the number of agents to 50 when we vary the number of rounds.

$$\text{Nash Welfare} = \sum_i e_i \log\left(\sum_r \min(d_{i,r}, a_{i,r})\right).$$

Observe that the Nash welfare metric is higher for the former scenario than the latter, which is in line with our intuition about which allocation is more fair.




Fig. 6. Sorted Sharing Index for Google Cluster Traces.




Fig. 7. Sorted Sharing Index for a Random Trace.

7.1 Performance Evaluation

Figure 3 presents social welfare from varied allocation mechanisms for both Google and random traces normalized to social welfare of static allocations. DMM and SMM produce the same, highest social welfare as they always allocate resources to those agents with high valuations. Note that SMM and DMM both fail to guarantee strategy-proofness when $L > 0$. Therefore, when agents report strategically, for all the mechanism knows, SMM and DMM's allocations could be as inefficient as static allocations. But this is not captured in the figure, which implicitly assumes truthful reporting.

The 1-Period mechanism produces the lowest social welfare. Increasing the period length to 2 slightly improves the welfare of the T -Period mechanism. Note that both mechanisms outperform static allocations. The $R/2$ -Period mechanism achieves 87% of SMM welfare for Google traces, but fails to provide strategy-proofness.

The social welfare of FL is competitive with state-of-the-art dynamic allocation mechanisms. FL achieves 97% of SMM's welfare for Google traces and 98% for random demands. In practice, strong game-theoretic desiderata do not come with high welfare costs.

Figure 4 compares the normalized Nash welfare from varied mechanisms. Once again, DMM and SMM outperform other mechanisms, but DMM and SMM's outcomes are no longer equal because the number of high-valued resources that each agent receives differs across mechanisms. FL achieves 99.7% of DMM welfare for both Google cluster and random traces. This high Nash welfare could be explained by FL's high social welfare and the fact that FL allocates agents their exact endowment across rounds.

Figure 5 shows how social welfare changes when varying population size and number of rounds under FL. As the population size increases, the diversity between agents' demands at each round increases. Agents' complementary demands improve welfare from FL as fewer agents are forced to spend tokens on low-valued resources. Moreover, as the number of rounds increases, agents' flexibility in spending their tokens on high-valued resources increases. We prove in §5.5 that, at least when endowments are equal, FL approximates efficiency.

7.2 Sharing Incentives

We define the *sharing index* of agent i to be the ratio between the number of high-valued resources agent i receives under FL and under static allocations. In §5.4, we show that FL guarantees that the sharing index of each agent is always at least 0.5. In practice, however, our simulations show that the sharing index is much higher.

Figure 6 shows the sharing index for all agents in the Google cluster traces, sorted in increasing order and shown on a log scale. The minimum sharing index across all agents is 0.98, and on average agents receive 15x more utility under FL compared to static allocations. As can be seen, there is high variance in sharing index across agents. Agents with high index are those who have zero demand at most of the rounds and very high demand at a few rounds. These agents benefit the most from sharing. When they have zero demand, they do not spend any tokens. Once they have a high demand they spend their tokens to receive the resources they need.

Figure 7 shows agents' sharing index for an instance with random demands. Since agents do not have correlated demands, the variance in sharing index is significantly lower compared to the Google cluster traces. Moreover, across all agents over 100 random instances, we do not observe a single violation of SI (*i.e.* no agent has a sharing index less than 1)

8 RELATED WORK

There is a body of work in the *mechanism design without money* literature that is related to our work. Gorokh et al. [18] consider a setting where a single item is to be allocated repeatedly, and extend to more general settings in a follow-up paper [19]. They do so by endowing each user with a fixed amount of artificial currency and then treating it similarly to if it were real money. They show that, for a large enough number of rounds, incentives to misreport and welfare loss both vanish. However, their notion of strategy-proofness is ex-ante Bayesian, requiring users (and the mechanism) to know the distribution from which other users' demands are drawn and truthful reporting is optimal only in expectation. Our notion of SP is ex-post, meaning that an agent never regrets truthful reporting.

Various other work does not explicitly use artificial currency, but by keeping track of how much utility an agent should receive in the future, achieve guarantees in a way that resembles the use of artificial currency [5, 9, 21]. Again, these results are for a weaker notion of SP.

In a similar setting, Aleksandrov et al. [6, 7] consider a stream of resources arriving one at a time that must be allocated among competing strategic agents. They consider two mechanisms, one of which is similar to SMM and the other similar to DMM. They obtain both positive and negative results for these mechanisms, however their positive results are primarily obtained for the case where agent utilities are 0 or 1, corresponding to our $L = 0$ case. They also consider only the symmetric agent setting, rather than our setting that allows unequal endowments.

There also exists literature on *dynamic fair division* [14, 23, 34], but this work predominantly focuses on agents arriving and departing over time, rather than the preferences themselves being dynamic, as in our work.

In the systems literature, in recent years, there has been a growing body of work on using economic game-theory to allocate resources [16, 35, 36]. These works only consider one-shot allocations and do not study allocations over time. Ghodsi et al. [15] consider dynamic allocations over time but in a completely different allocation setting than ours. Their proposed mechanism allocates resources to packets in a queue. In such a setting, time cannot be divided into fixed intervals, because processing packets take different times, which means a packet could stall all other packets until its processed. As a result, proportional allocations have to be approximated through discrete packet scheduling decisions [13, 26].

In a work that is close to our setting, Tang et al. [30] propose a dynamic allocation policy that resembles DMM. We study the characteristics of DMM in §3 and evaluate its performance in §7. Another related work in this area is that of Sandholm and Lai [28]. The authors propose a scheduler that allocates resources between users with dynamically changing demands. This work deploys heuristics and does not provide any theoretical guarantees that we study in this paper.

9 CONCLUSION

We have considered the problem of designing mechanisms for dynamic proportional sharing in a high-low utility model that both incentivize users to participate and share their resources (sharing incentives), as well as truthfully report their resource requirements to the system (strategy-proofness). We show that while each of these properties is incompatible with full efficiency, it is possible to satisfy both of them and still obtain some efficiency gains from sharing.

The main mechanism that we present, the flexible lending mechanism, is strategy-proof and provides each user a theoretical guarantee of at least half her sharing incentives share. While we do not guarantee full sharing incentives, we show via simulations on both real and synthetic data that in practical situations, no users are significantly worse off by participating in the sharing scheme (and the majority are vastly better off). We show that under certain assumptions, the flexible lending mechanism provides full efficiency in the large round limit, which is supported by our simulation results. By incentivizing truthful reporting, we posit that the flexible lending mechanism will in fact produce significant efficiency gains in settings where agents are strategic.

Many directions for future work remain. The 2-Period mechanism fully satisfies both SP and SI, but remains very inflexible in its allocations. A key challenge is the design of a more flexible mechanism that satisfies both properties (or some upper bound on the efficiency that such mechanisms can achieve). Another direction is to extend the utility model. The high/low model is crucial to the positive strategic results that we obtain because trade-offs are well-defined: swapping an L resource for an H resource is always bad. Even introducing a medium (M) value complicates the situation considerably, and extending to such a setting would represent an exciting step forward.

ACKNOWLEDGMENTS

This work is supported by the National Science Foundation under grants CCF-1149252 (CAREER), CCF-1337215 (XPS-CLCCA), SHF-1527610, AF-1408784, CCF-133721 and IIS-1527434. This work is also supported by STARnet, a Semiconductor Research Corporation program, sponsored by MARCO and DARPA. Freeman is supported by a Facebook graduate fellowship. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of these sponsors.

REFERENCES

- [1] 2011. More Google cluster data. Google research blog. (Nov. 2011). Posted at <http://googleresearch.blogspot.com/2011/11/more-google-cluster-data.html>.
- [2] Accessed Jan 2018. Capacity Scheduler. <https://hadoop.apache.org/docs/current/hadoop-yarn/hadoop-yarn-site/CapacityScheduler.html>. (Accessed Jan 2018).
- [3] Accessed Jan 2018. Fair Scheduler. <https://hadoop.apache.org/docs/current/hadoop-yarn/hadoop-yarn-site/FairScheduler.html>. (Accessed Jan 2018).
- [4] Accessed Jan 2018. Spark Dynamic Allocation. <https://spark.apache.org/docs/latest/job-scheduling.html>. (Accessed Jan 2018).
- [5] Atila Abdulkadiroğlu and Kyle Bagwell. 2013. Trust, Reciprocity, and Favors in Cooperative Relationships. *American Economic Journal: Microeconomics* 5, 2 (2013), 213–259.
- [6] Martin Aleksandrov, Haris Aziz, Serge Gaspers, and Toby Walsh. 2015. Online Fair Division: Analysing a Food Bank Problem.. In *IJCAI*. 2540–2546.

- [7] Martin Aleksandrov and Toby Walsh. 2017. Pure Nash Equilibria in Online Fair Division. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. 42–48.
- [8] Anonymous. 2017. SHARCNET. <https://www.sharcnet.ca>. (2017).
- [9] Susan Athey and Kyle Bagwell. 2001. Optimal Collusion with Private Information. *RAND Journal of Economics* 32, 3 (2001), 428–465.
- [10] Luiz André Barroso and Urs Hölzle. 2007. The case for energy-proportional computing. *Computer* 40, 12 (2007).
- [11] Garrett Birkhoff. 1946. Tres observaciones sobre el algebra lineal. *Univ. Nac. Tucumán Rev. Ser. A* 5 (1946), 147–151.
- [12] Eric Boutin, Jaliya Ekanayake, Wei Lin, Bing Shi, Jingren Zhou, Zhengping Qian, Ming Wu, and Lidong Zhou. 2014. Apollo: Scalable and coordinated scheduling for cloud-scale computing. In *Proceedings of the 11th USENIX conference on Operating Systems Design and Implementation (OSDI)*. USENIX Association, 285–300.
- [13] Alan Demers, Srinivasan Keshav, and Scott Shenker. 1989. Analysis and simulation of a fair queueing algorithm. In *ACM SIGCOMM Computer Communication Review*, Vol. 19. ACM, 1–12.
- [14] Eric Friedman, Christos-Alexandros Psomas, and Shai Vardi. 2017. Controlled Dynamic Fair Division. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. ACM, 461–478.
- [15] Ali Ghodsi, Vyas Sekar, Matei Zaharia, and Ion Stoica. 2012. Multi-resource fair queueing for packet processing. *ACM SIGCOMM Computer Communication Review* 42, 4 (2012), 1–12.
- [16] Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. 2011. Dominant Resource Fairness: Fair Allocation of Multiple Resource Types.. In *NSDI*, Vol. 11. 24–24.
- [17] Ali Ghodsi, Matei Zaharia, Scott Shenker, and Ion Stoica. 2013. Choosy: Max-min fair sharing for datacenter jobs with constraints. In *Proceedings of the 8th ACM European Conference on Computer Systems*. ACM, 365–378.
- [18] Artur Gorokh, Siddhartha Banerjee, and Krishnamurthy Iyer. 2016. Near-Efficient Allocation Using Artificial Currency in Repeated Settings. (2016).
- [19] Artur Gorokh, Siddhartha Banerjee, and Krishnamurthy Iyer. 2017. From Monetary to Non-Monetary Mechanism Design Via Artificial Currencies. (2017).
- [20] Ajay Gulati, Ganesha Shanmuganathan, Xuechen Zhang, and Peter J Varman. 2012. Demand Based Hierarchical QoS Using Storage Resource Pools.. In *USENIX Annual Technical Conference*. 1–13.
- [21] Mingyu Guo, Vincent Conitzer, and Daniel M. Reeves. 2009. Competitive Repeated Allocation Without Payments. In *Proceedings of the Fifth Workshop on Internet and Network Economics (WINE)*. Rome, Italy, 244–255.
- [22] Gary J Henry. 1984. The UNIX system: The fair share scheduler. *AT&T Bell Laboratories Technical Journal* 63, 8 (1984), 1845–1857.
- [23] Ian Kash, Ariel D Procaccia, and Nisarg Shah. 2014. No agent left behind: Dynamic fair division of multiple resources. *Journal of Artificial Intelligence Research* 51 (2014), 579–603.
- [24] J. Kay and P. Lauder. 1988. A Fair Share Scheduler. *Commun. ACM* 31, 1 (1988), 44–55.
- [25] Lawrence Berkeley National Laboratory. 2017. National Energy Research Scientific Computing Center. <https://www.nersc.gov>. (2017).
- [26] Abhay K Parekh and Robert G Gallager. 1993. A generalized processor sharing approach to flow control in integrated services networks: The single-node case. *IEEE/ACM Transactions on Networking* 1, 3 (1993), 344–357.
- [27] Charles Reiss, John Wilkes, and Joseph L. Hellerstein. 2011. *Google cluster-usage traces: format + schema*. Technical Report. Google Inc., Mountain View, CA, USA. Revised 2014-11-17 for version 2.1. Posted at <https://github.com/google/cluster-data>.
- [28] Thomas Sandholm and Kevin Lai. 2010. Dynamic proportional share scheduling in Hadoop. In *Job scheduling strategies for parallel processing*. Springer, 110–131.
- [29] Alan Shieh, Srikanth Kandula, Albert G Greenberg, Changhoon Kim, and Bikas Saha. 2011. Sharing the Data Center Network.. In *NSDI*, Vol. 11. 23–23.
- [30] Shanjiang Tang, Bu-Sung Lee, Bingsheng He, and Haikun Liu. 2014. Long-term resource fairness: Towards economic fairness on pay-as-you-use computing systems. In *Proceedings of the 28th ACM international conference on Supercomputing*. ACM, 251–260.
- [31] Duke University. 2017. The Duke Computer Cluster. <http://rc.duke.edu/the-duke-compute-cluster.edu>. (2017).
- [32] Abhishek Verma, Luis Pedrosa, Madhukar Korupolu, David Oppenheimer, Eric Tune, and John Wilkes. 2015. Large-scale Cluster Management at Google with Borg. In *Proceedings of the Tenth European Conference on Computer Systems (EuroSys)*. ACM, 18:1–18:17.
- [33] Carl A. Waldspurger and William E. Weihl. 1994. Lottery Scheduling: Flexible Proportional-Share Resource Management. In *Proceedings of the First Symposium on Operating Systems Design and Implementation (OSDI)*. 1–11.
- [34] Toby Walsh. 2011. Online cake cutting. *Algorithmic Decision Theory* (2011), 292–305.
- [35] Xiaodong Wang and José F Martínez. 2015. XChange: Scalable Dynamic Multi-Resource Allocation in Multicore Architectures. In *Proceedings of the 21st IEEE International Symposium on High Performance Computer Architecture (HPCA)*.

[36] Seyed Majid Zahedi and Benjamin C. Lee. 2014. REF: Resource Elasticity Fairness with Sharing Incentives for Multiprocessors. In *Proceedings of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. ACM, 145–160.

A OMITTED PROOFS

A.1 Proof of Lemma 14

We use the characterization of the FL mechanism allocations from Lemma 13. We consider four cases, corresponding to whether or not supply exceeds demand in the truthful and misreported instances. Let x' denote the objective value in the FL mechanism's call to PSWC in the misreported instance, and x in the truthful instance. Suppose first that $D_r \geq E$ and $D'_r \geq E$. Suppose that $x' \leq x$. Then, for all $j \neq i$,

$$a'_{j,r} = \min(x'e_j, d_{j,r}, t'_{j,r}) \leq \min(xe_j, d_{j,r}, t_{j,r}) = a_{j,r},$$

which implies that $a'_{i,r} \geq a_{i,r}$, since $\sum_{k \in [n]} a_{k,r} = \sum_{k \in [n]} a'_{k,r}$. On the other hand, if $x' > x$, then

$$a'_{i,r} = \min(x'e_i, d_{i,r}, t'_{i,r}) \geq \min(xe_i, d_{i,r}, t_{i,r}) = a_{i,r}.$$

Second, suppose that $D_r \geq E$ and $D'_r < E$. Then

$$a'_{i,r} \geq \min(d_{i,r}, t'_{i,r}) \geq \min(d_{i,r}, t_{i,r}) \geq a_{i,r}.$$

Third, suppose that $D_r < E$ and $D'_r \geq E$. Then

$$a'_{j,r} \leq \min(d_{j,r}, t'_{j,r}) \leq \min(d_{j,r}, t_{j,r}) \leq a_{j,r}$$

for all $j \neq i$, which implies that $a'_{i,r} \geq a_{i,r}$. Finally, suppose that $D_r < E$ and $D'_r < E$. If $x' \leq x$, then for all $j \neq i$, we have that

$$a'_{j,r} = \min(t'_{j,r}, \max(d_{j,r}, x'e_j)) \leq \min(t_{j,r}, \max(d_{j,r}, xe_j)) = a_{j,r},$$

which implies that $a'_{i,r} \geq a_{i,r}$. If $x' > x$, then

$$a'_{i,r} = \min(t'_{i,r}, \max(d_{i,r}, x'e_i)) \geq \min(t_{i,r}, \max(d_{i,r}, xe_i)) = a_{i,r}.$$

Thus, the lemma holds in all cases.

A.2 Proof of Lemma 27

If $D \geq E$, substituting the relevant terms into Lemma 11 gives us the following.

$$a_{i,r} = \max(0, \min(\min(d'_{i,r}, e_i + b_i), xe_i)) = \min(e_i + b_i, d'_{i,r}, xe_i).$$

If $D < E$, then again by substituting into Lemma 11 we have the following.

$$a_{i,r} = \max(\min(e_i + b_i, d'_{i,r}), \min(e_i + b_i, xe_i)) = \min(e_i + b_i, \max(d'_{i,r}, xe_i)).$$

The final equality, $\max(\min(A, B), \min(A, C)) = \min(A, \max(B, C))$, can easily be checked to hold case by case for any relative ordering of A , B , and C .

A.3 Proof of Lemma 28

If $r > T$, then the allocation of agent i is independent of her reported demand, thus $a_{i,r} = a'_{i,r}$. Now suppose that $r \leq T$. Let $\bar{d}_{i,r} = \min(d_{i,r}, e_i + b_{i,r})$ and $\bar{d}'_{i,r} = \min(d'_{i,r}, e_i + b_{i,r})$. Also, let x and x' denote the objective value in the T -period mechanism's call to PSWC when i reports $d_{i,r}$ and $d'_{i,r}$, respectively. Observe first that $D' = \bar{d}'_{i,r} + \sum_{j \neq i} \bar{d}_{j,r} \leq \bar{d}_{i,r} + \sum_{j \neq i} \bar{d}_{j,r} = D$.

Suppose first that $E \leq D' \leq D$. Let $a_{j,r}$ and $a'_{j,r}$ denote the allocations of j when i reports $d_{i,r}$ and $d'_{i,r}$, respectively. If $x' \geq x$, then for all $j \neq i$, by Lemma 27 we have:

$$a'_{j,r} = \min(e_j + b_{j,r}, d_{j,r}, x'e_j) \geq \min(e_j + b_{j,r}, d_{j,r}, xe_j) = a_{j,r}.$$

This immediately implies that $a'_{i,r} \leq a_{i,r}$, because $\sum_{k \in [n]} a_{k,r} = \sum_{k \in [n]} a'_{k,r} = E$. If $x' < x$, then again by Lemma 27 we have the following:

$$a'_{i,r} = \min(e_i + b_{i,r}, d'_{i,r}, x'e_i) \leq \min(e_i + b_{i,r}, d_{i,r}, xe_i) = a_{i,r}.$$

By the same lemma, for all $j \neq i$, we also have:

$$a'_{j,r} = \min(e_j + b_{j,r}, d_{j,r}, x'e_j) \leq \min(e_j + b_{j,r}, d_{j,r}, xe_j) = a_{j,r}.$$

Therefore, for all $k \in [n]$, $a_{k,r} \geq a'_{k,r}$. However, since $\sum_{k \in [n]} a_{k,r} = \sum_{k \in [n]} a'_{k,r} = E$, it has to be the case that $a_{k,r} = a'_{k,r}$ for all k .

Next, suppose that $D' < E \leq D$. By the definition of the T -period mechanism, for all $j \neq i$, $a'_{j,r} \geq \bar{d}_{j,r}$, and $a_{j,r} \leq \bar{d}_{j,r}$. Therefore, $a'_{j,r} \geq a_{j,r}$ which implies that $a'_{i,r} \leq a_{i,r}$.

Finally, suppose that $D' \leq D < E$. If $x' \geq x$, then by Lemma 27, for all $j \neq i$, we have:

$$a'_{j,r} = \min(e_j + b_{j,r}, \max(d_{j,r}, x'e_j)) \geq \min(e_j + b_{j,r}, \max(d_{j,r}, xe_j)) = a_{j,r}.$$

This implies $a'_{i,r} \leq a_{i,r}$. If $x' < x$, then, by Lemma 27 we have:

$$a'_{i,r} = \min(e_i + b_{i,r}, \max(d'_{i,r}, x'e_i)) \leq \min(e_i + b_{i,r}, \max(d_{i,r}, xe_i)) = a_{i,r}$$

By the same lemma, for all $j \neq i$, we also have:

$$a'_{j,r} = \min(e_j + b_{j,r}, \max(d_{j,r}, x'e_j)) \leq \min(e_j + b_{j,r}, \max(d_{j,r}, xe_j)) = a_{j,r}.$$

Therefore, for all $k \in [n]$, $a'_{k,r} \leq a_{k,r}$. However, since $\sum_{k \in [n]} a_{k,r} = \sum_{k \in [n]} a'_{k,r} = E$, it has to be the case that $a_{k,r} = a'_{k,r}$ for all k .

A.4 Proof of Lemma 29

Note that $D' \leq D$, since $d'_{i,r} < d_{i,r}$ and $d'_{j,r} = d_{j,r}$ for all agents $j \neq i$. If $E \leq D$, then by the definition of the T -period mechanism we have:

$$a_{i,r} \leq \bar{d}_{i,r} = \min(e_i + b_{i,r}, d_{i,r}) \leq d_{i,r}.$$

Next, assume that $D' \leq D < E$. Then $a'_{i,r} < a_{i,r}$ implies that there is at least one agent j with $a'_{j,r} > a_{j,r}$. In the proof of Lemma 28 we show that if $x' < x$, then $a'_{k,r} = a_{k,r}$ for all k . Therefore, it has to be the case that $x' \geq x$. By Lemma 27, $a_{i,r} = \min(e_i + b_{i,r}, \max(d_{i,r}, xe_i))$ and $a'_{i,r} = \min(e_i + b_{i,r}, \max(d'_{i,r}, x'e_i))$. It is easy to see that if $d_{i,r} < xe_i$, then $a'_{i,r} \geq a_{i,r}$, which contradicts the assumption in the lemma statement. Therefore, we have:

$$a_{i,r} = \min(e_i + b_{i,r}, \max(d_{i,r}, xe_i)) = \min(e_i + b_{i,r}, d_{i,r}) \leq d_{i,r}.$$

A.5 Proof of Corollary 30

Because $a'_{i,r} < a_{i,r} \leq d_{i,r}$, we can substitute the utility values from Equation (2),

$$u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r}) = a_{i,r}H - a'_{i,r}H = H(a_{i,r} - a'_{i,r}).$$

A.6 Proof of Lemma 31

Note that $D' \geq D$, since $d'_{i,r} > d_{i,r}$ and $d'_{j,r} = d_{j,r}$ for all agents $j \neq i$. If $D < E$, then $a_{i,r} \geq \bar{d}_{i,r} = \min(e_i + b_{i,r}, d_{i,r})$. We show that $e_i + b_{i,r} \geq d_{i,r}$, and therefore, $a_{i,r} \geq d_{i,r}$. Suppose for contradiction that $e_i + b_{i,r} < d_{i,r} < d'_{i,r}$, which means $\bar{d}_{i,r} = \bar{d}'_{i,r} = e_i + b_{i,r}$. By definition of the T -period mechanism, $\bar{d}_{i,r} \leq a_{i,r} \leq e_i + b_{i,r}$, which implies $a_{i,r} = e_i + b_{i,r}$. Also, by the definition of the mechanism, $a'_{i,r} \leq \bar{d}'_{i,r} = e_i + b_{i,r}$ if $D' \geq E$, and $a'_{i,r} \leq e_i + b'_{i,r} = e_i + b_{i,r}$ if $D' < E$. In both cases, $a'_{i,r} \leq e_i + b_{i,r} = a_{i,r}$, a contradiction to the assumption in the lemma statement.

If $D' \geq D \geq E$, then $a'_{i,r} > a_{i,r}$ implies that there is at least an agent j with $a'_{j,r} < a_{j,r}$. In the proof of Lemma 28 we show that if $x < x'$, then $a'_{k,r} = a_{k,r}$ for all k . Therefore, it has to be the

case that $x \geq x'$. By Lemma 27, $a_{i,r} = \min(e_i + b_{i,r}, d_{i,r}, xe_i)$ and $a'_{i,r} = \min(e_i + b_{i,r}, d'_{i,r}, x'e_i)$. It is easy to see that if $a_{i,r}$ is $x'e_i$ or $e_i + b_{i,r}$, then $a'_{i,r} \leq a_{i,r}$. Therefore, $a_{i,r} = d_{i,r}$, which means the lemma holds.

A.7 Proof of Corollary 32

Because $d_{i,r} \leq a_{i,r} < a'_{i,r}$, we can substitute the utility values from Equation (2),

$$u_{i,r}(a'_{i,r}) - u_{i,r}(a_{i,r}) = d_{i,r}H + (a'_{i,r} - d_{i,r})L - (d_{i,r}H + (a_{i,r} - d_{i,r})L) = L(a'_{i,r} - a_{i,r}).$$

A.8 Proof of Lemma 33

Suppose first that agent i reports $d'_{i,T} < d_{i,T}$. Then, by Lemma 28, $a'_{i,T} \leq a_{i,T}$. If $a'_{i,T} = a_{i,T}$, then the misreport has had no effect on the allocations, since the allocation at rounds $r \leq T$ is unchanged, and the allocations at rounds $r > T$ depend only on the allocations at rounds $r \leq T$, not the reports. So assume that $a'_{i,T} = a_{i,T} - k$ for some $k > 0$. By the definition of the T -Period mechanism, i 's allocation increases by $\frac{k}{T}$ for each of rounds $T + 1, \dots, 2T$. The difference between her utility from truthfully reporting at round T and from misreporting at round T is given by

$$\begin{aligned} U_{i,R}(\mathbf{a}_i) - U_{i,R}(\mathbf{a}'_i) &= \sum_{r=1}^R (u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r})) \\ &= u_{i,T}(a_{i,T}) - u_{i,T}(a'_{i,T}) + \sum_{r=T+1}^{2T} (u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r})) \\ &= kH + \sum_{r=T+1}^{2T} \left(u_{i,r}(a_{i,r}) - u_{i,r}\left(a_{i,r} + \frac{k}{T}\right) \right) \geq kH - kH = 0 \end{aligned}$$

where the second transition follows because $d'_{i,r} = d_{i,r}$ for all rounds $r < T$, the third transition from Corollary 30, and the final transition because each of the extra resources received in the misreported case for rounds $r > T$ can each be worth at most H to i .

Next suppose that agent i reports $d'_{i,T} > d_{i,T}$. Then, by Lemma 28, $a'_{i,T} \geq a_{i,T}$. As before, assume that $a'_{i,T} \neq a_{i,T}$. That is, $a'_{i,T} = a_{i,T} + k$ for some $k > 0$. By the definition of the T -Period mechanism, i 's allocation decreases by $\frac{k}{T}$ for each of rounds $T + 1, \dots, 2T$. The difference between her utility from truthfully reporting at round T and from misreporting at round T is given by

$$\begin{aligned} U_{i,R}(\mathbf{a}_i) - U_{i,R}(\mathbf{a}'_i) &= \sum_{r=1}^R (u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r})) \\ &= u_{i,T}(a_{i,T}) - u_{i,T}(a'_{i,T}) + \sum_{r=T+1}^{2T} (u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r})) \\ &= -kL + \sum_{r=T+1}^{2T} \left(u_{i,r}(a_{i,r}) - u_{i,r}\left(a_{i,r} - \frac{k}{T}\right) \right) \geq -kL + kL = 0 \end{aligned}$$

where the second transition follows because $d'_{i,r} = d_{i,r}$ for all rounds $r < T$, the third transition from Corollary 32, and the final transition because each of the extra resources received in the truthful case for rounds $r > T$ are each worth at least L to i .

A.9 Proof of Lemma 35

We treat four cases, corresponding to whether or not supply exceeds demand in the truthful and misreported instances. Let x' denote the objective value in the T -Period mechanism's call to

PSWC in the misreported instance, and x in the truthful instance. All cases rely heavily on the characterization of the allocation from Lemma 27.

Suppose first that $D_r \geq E$ and $D'_r \geq E$. Suppose that $x' \leq x$. Then, for all $j \neq i$, $a'_{j,r} = \min(x'e_j, d_{j,r}, e_j + b'_{j,r}) \leq \min(xe_j, d_{j,r}, e_j + b_{j,r}) = a_{j,r}$, which implies that $a'_{i,r} \geq a_{i,r}$, since $\sum_{k \in [n]} a_{k,r'} = \sum_{k \in [n]} a'_{k,r'}$. On the other hand, if $x' > x$, then $a'_{i,r} = \min(x'e_i, d_{i,r}, e_i + b'_{i,r}) \geq \min(xe_i, d_{i,r}, e_i + b_{i,r}) = a_{i,r}$. Second, suppose that $D_r \geq E$ and $D'_r < E$. Then $a'_{i,r} \geq \min(d_{i,r}, e_i + b'_{i,r}) \geq \min(d_{i,r}, e_i + b_{i,r}) \geq a_{i,r}$. Third, suppose that $D_r < E$ and $D'_r \geq E$. Then $a'_{j,r} \leq \min(d_{j,r}, e_j + b'_{j,r}) \leq \min(d_{j,r}, e_j + b_{j,r}) \leq a_{j,r}$ for all $j \neq i$, which implies that $a'_{i,r} \geq a_{i,r}$.

Finally, suppose that $D_r < E$ and $D'_r < E$. If $x' \leq x$, then for all $j \neq i$, we have that $a'_{j,r} = \min(e_j + b'_{j,r}, \max(d_{j,r}, x'e_j)) \leq \min(e_j + b_{j,r}, \max(d_{j,r}, xe_j)) = a_{j,r}$, which implies that $a'_{i,r} \geq a_{i,r}$. If $x' > x$, then $a'_{i,r} = \min(e_i + b'_{i,r}, \max(d_{i,r}, x'e_i)) \geq \min(e_i + b_{i,r}, \max(d_{i,r}, xe_i)) = a_{i,r}$. Thus, the lemma holds in all cases.

A.10 Proof of Proposition 37

Let $r \leq T$. First suppose that $D < E$. Then i 's minimum allocation is $\bar{d}_{i,r} = \min(d'_{i,r}, e_i + b_{i,r}) = e_i$. So we know that $a_{i,r} \geq e_i$. Suppose for contradiction that $a_{i,r} > e_i$. Then there must be some agent $j \neq i$ with $a_{j,r} \leq e_j$. But now we could obtain a smaller value of x in the PSWC program by assigning slightly higher allocation to j , and slightly lower allocation to any agent with $a_{k,r}/e_k = x$ (we know that j is not one of these agents since $a_{j,r}/e_j < 1 < a_{i,r}/e_i \leq x$). This contradicts optimality of the PSWC program, therefore $a_{i,r} = e_i$.

Next, suppose that $D \geq E$. Then i 's limit allocation is $\bar{d}_{i,r} = \min(d'_{i,r}, e_i + b_{i,r}) = e_i$. So we know that $a_{i,r} \leq e_i$. Suppose for contradiction that $a_{i,r} < e_i$. Then there must exist some agent j with $a_{j,r} > e_j$. But now the objective value x of the call to PSWC could be improved by transferring some small amount of allocation to i from all agents k with $a_{k,r}/e_k = x$ (we know that i is not one of these agents since $a_{i,r}/e_i < 1 < a_{j,r}/e_j \leq x$). This contradicts optimality of the PSWC program, therefore $a_{i,r} = e_i$.

B OVER-REPORTING DEMAND IS NOT ADVANTAGEOUS

In this section we assume that $d'_{i,r'} > d_{i,r'}$. The setup otherwise mirrors that of §5.3.

LEMMA 40. *For all agents $j \neq i$, we have that $a'_{j,r'} \leq a_{j,r'}$. Further, $a'_{i,r'} \geq a_{i,r'}$.*

PROOF. We prove the statement for all $j \neq i$. The statement for i follows immediately because the total number of resources to allocate is fixed.

Observe first that

$$D_{r'} = \sum_{k \in [n]} \min(d_{k,r'}, t_{k,r'}) \leq \sum_{k \in [n]} \min(d'_{k,r'}, t_{k,r'}) = D'_{r'},$$

since i 's demand increases in the misreported instances but all other demands and token counts stay the same. Let x' denote the objective value in FL's call to PSWC in the misreported instance, and x in the truthful instance.

Suppose that $E \leq D_{r'} \leq D'_{r'}$. Suppose first that $x' < x$. Then, by Lemma 13,

$$a_{j,r'} = \min(xe_j, d_{j,r'}, t_{j,r'}) \geq \min(x'e_j, d_{j,r'}, t_{j,r'}) = a'_{j,r'}$$

for all $j \neq i$. Next, suppose that $x' \geq x$. Then, again by Lemma 13 and the fact that $d'_{i,r'} > d_{i,r'}$,

$$a'_{i,r'} = \min(x'e_i, d'_{i,r'}, t_{i,r'}) \geq \min(xe_i, d_{i,r'}, t_{i,r'}) = a_{i,r'}$$

And, for all $j \neq i$,

$$a'_{j,r'} = \min(x'e_j, d_{j,r'}, t_{j,r'}) \geq \min(xe_j, d_{j,r'}, t_{j,r'}) = a_{j,r'}$$

Because $a'_{k,r'} \geq a_{k,r'}$ for all users k , and $\sum_{k \in [n]} a_{k,r'} = \sum_{k \in [n]} a'_{k,r'}$, it must be the case that $a'_{k,r'} = a_{k,r'}$ for all k , which satisfies the statement of the lemma.

Next, suppose that $D_{r'} < E \leq D'_{r'}$. By the definition of FL, $a_{k,r'} \geq \min(d_{k,r'}, t_{k,r'})$ for all k , and $a'_{k,r'} \leq \min(d'_{k,r'}, t_{k,r'})$ for all k . Since $\min(d'_{j,r'}, t_{j,r'}) = \min(d_{j,r'}, t_{j,r'})$ for all $j \neq i$, we have that $a_{j,r'} \geq a'_{j,r'}$, implying also that $a_{i,r'} \leq a'_{i,r'}$.

Finally, suppose that $D_{r'} \leq D'_{r'} < E$. Suppose first that $x \leq x'$. Then, by Lemma 13 and the assumption that $d_{i,r'} < d'_{i,r'}$, we have

$$a_{i,r'} = \min(t_{i,r'}, \max(xe_i, d_{i,r'})) \leq \min(t_{i,r'}, \max(x'e_i, d'_{i,r'})) = a'_{i,r'}$$

and

$$a_{j,r'} = \min(t_{j,r'}, \max(xe_j, d_{j,r'})) \leq \min(t_{j,r'}, \max(x'e_j, d_{j,r'})) = a'_{j,r'}$$

for all $j \neq i$. Because $a_{k,r'} \leq a'_{k,r'}$ for all users k , and $\sum_{k \in [n]} a'_{k,r'} = \sum_{k \in [n]} a_{k,r'}$, it must be the case that $a_{k,r'} = a'_{k,r'}$ for all k , which satisfies the lemma statement. Next, suppose that $x > x'$. Then, again by Lemma 13, for all $j \neq i$, we have

$$a_{j,r'} = \min(t_{j,r'}, \max(xe_j, d_{j,r'})) \geq \min(t_{j,r'}, \max(x'e_j, d_{j,r'})) = a'_{j,r'}.$$

□

If it is the case that $a'_{i,r'} = a_{i,r'}$, then it must also be that $a'_{j,r'} = a_{j,r'}$ for all $j \neq i$. So allocations at round r' are the same in the misreported instance as the truthful instance. Therefore, for all rounds $r \leq r'$, allocations in both universes are the same. In all rounds $r > r'$, reports in both universes are the same. Together, these imply that allocations for all rounds $r > r'$ are the same in both universes. In particular, i does not profit from her misreport and could weakly improve her utility by reporting $d'_{i,r'} = d_{i,r'}$. So, for the remainder of this section, we assume that $a'_{i,r'} > a_{i,r'}$.

Our next lemma says that the additional resources that i receives in round r' are *low* valued resources for her. The intuition is that if it were the case that i was receiving only high-valued resources under truthful reporting, then she will not receive any extra resources by misreporting (since no agent donates any additional resources for i to receive).

LEMMA 41. *If $a'_{i,r'} > a_{i,r'}$, then $a_{i,r'} \geq d_{i,r'}$.*

PROOF. Suppose for contradiction that $a_{i,r'} < d_{i,r'}$. We also know that $a_{i,r'} < a'_{i,r'} \leq t'_{i,r'} = t_{i,r'}$, where the equality holds because allocations before round r' are identical in the truthful and misreported instances. It must therefore be the case that $D'_{r'} \geq D_{r'} > E$, where the first inequality holds because $d'_{j,r'} = d_{j,r'}$ for all $j \neq i$ and $d'_{i,r'} > d_{i,r'}$, and the second because $a_{i,r'} < \min(t_{i,r'}, d_{i,r'})$. Let x denote the objective value of FL's call to PSWC in the truthful instance, and x' in the misreported instance. Suppose that $x \leq x'$. Then, by Lemma 13 and the assumption that $d_{i,r'} < d'_{i,r'}$,

$$a_{i,r'} = \min(t_{i,r'}, xe_i, d_{i,r'}) \leq \min(t_{i,r'}, x'e_i, d'_{i,r'}) = a'_{i,r'},$$

and for all $j \neq i$

$$a_{j,r'} = \min(t_{j,r'}, xe_j, d_{j,r'}) \leq \min(t_{j,r'}, x'e_j, d_{j,r'}) = a'_{j,r'}.$$

Because $a_{k,r'} \leq a'_{k,r'}$ for all agents k , and $\sum_{k \in [n]} a_{k,r'} = \sum_{k \in [n]} a'_{k,r'}$, it must be the case that $a'_{k,r'} = a_{k,r'}$ for all k . This contradicts the assumption that $a_{i,r'} < a'_{i,r'}$.

Now suppose that $x > x'$. Note that $xe_i < d_{i,r'} < d'_{i,r'}$, where the first inequality holds because $a_{i,r'} < \min(t_{i,r'}, d_{i,r'})$. Then, again by Lemma 13 and the previous observation, we have

$$a'_{i,r'} = \min(t_{i,r'}, x'e_i, d'_{i,r'}) \leq \min(t_{i,r'}, xe_i, d_{i,r'}) = a_{i,r'},$$

which contradicts that $a_{i,r'} < a'_{i,r'}$.

Since we arrive at a contradiction in all cases, the lemma statement must be true. □

As a corollary, we can write the difference in utility between the truthful and misreported instances that i derives from round r' .

$$\text{COROLLARY 42. } u_{i,r'}(a'_{i,r'}) - u_{i,r'}(a_{i,r'}) = L(a'_{i,r'} - a_{i,r'}).$$

PROOF. Because $d_{i,r'} \leq a_{i,r'} < a'_{i,r'}$, we can substitute the utility values from Equation (2):

$$u_{i,r'}(a'_{i,r'}) - u_{i,r'}(a_{i,r'}) = d_{i,r'}H + (a'_{i,r'} - d_{i,r'})L - d_{i,r'}H - (a'_{i,r'} - d_{i,r'})L = L(a'_{i,r'} - a_{i,r'}).$$

□

For a fixed agent k , denote by r'_k the round at which agent k runs out of tokens in the misreported universe. That is, r'_k is the first (and only) round with $a'_{k,r'_k} = t'_{k,r'_k} > 0$. Note that $r'_i \geq r'$, since $a'_{i,r'} > 0$. Given this, our next lemma states that, under certain conditions, the effect of i 's misreport, $d'_{i,r} > d_{i,r}$, is to decrease the objective value of FL's call to PSWC.

LEMMA 43. *Let $r < r'_i$ (i.e. $a'_{i,r} < t'_{i,r}$). Suppose $t_{j,r} \leq t'_{j,r}$ for all agents $j \neq i$. Suppose that either $\min(D_r, D'_r) \geq E$ or $\max(D_r, D'_r) < E$. Then $x' \leq x$, where x' denotes the objective value of FL's call to PSWC in the misreported instance and x in the truthful instance.*

PROOF. First, suppose that $\min(D_r, D'_r) \geq E$. Suppose for contradiction that $x < x'$. By Lemma 13,

$$a_{j,r} = \min(xe_j, d_{j,r}, t_{j,r}) \leq \min(x'e_j, d_{j,r}, t'_{j,r}) = a'_{j,r}$$

for all $j \neq i$, where the inequality follows from the assumption that $x < x'$ and that $t_{j,r} \leq t'_{j,r}$. Further,

$$\begin{aligned} a_{i,r} &= \min(xe_i, d_{i,r}, t_{i,r}) \leq \min(xe_i, d_{i,r}) \leq \min(x'e_i, d_{i,r}) \\ &= \min(x'e_i, d_{i,r}, t'_{i,r}) = a'_{i,r}, \end{aligned}$$

where the second inequality follows from the assumption that $x < x'$ and the second to last equality from the assumption $a'_{i,r} < t'_{i,r}$.

Therefore, $a_{k,r} \leq a'_{k,r}$ for all agents k . Since $\sum a'_{k,r} = \sum a_{k,r}$, it must be the case that $a'_{k,r} = a_{k,r}$ for all agents k . Therefore, by the definition of FL, $a'_{k,r}/e_k \leq x < x'$ for all agents k with $a'_{k,r} > m_k = 0$. Therefore x' is not the optimal objective value of the PSWC program in the misreported instance, a contradiction. Thus, $x \geq x'$.

Next, suppose that $\max(D_r, D'_r) < E$. Suppose for contradiction that $x < x'$. By Lemma 13,

$$a_{j,r} = \min(t_{j,r}, \max(xe_j, d_{j,r})) \leq \min(t'_{j,r}, \max(x'e_j, d_{j,r})) = a'_{j,r}$$

for all $j \neq i$, where the inequality follows from the assumption that $x < x'$ and that $t_{j,r} \leq t'_{j,r}$. Further,

$$\begin{aligned} a_{i,r} &= \min(t_{i,r}, \max(xe_i, d_{i,r})) \leq \max(xe_i, d_{i,r}) \leq \max(x'e_i, d_{i,r}) \\ &= \min(t'_{i,r}, \max(x'e_i, d_{i,r})) = a'_{i,r}, \end{aligned}$$

where the second inequality follows from the assumption that $x < x'$ and the second to last equality from the assumption $a'_{i,r} < t'_{i,r}$.

Therefore, $a_{k,r} \leq a'_{k,r}$ for all agents k . Since $\sum a'_{k,r} = \sum a_{k,r}$, it must be the case that $a'_{k,r} = a_{k,r}$ for all agents k . Consider all agents with $\min(d_{k,r}, t'_{k,r}) < a'_{k,r}$ (that is, those agents for which the first constraint in the PSWC program binds in the misreported instance). For all such agents, we have

$$\min(d_{k,r}, t'_{k,r}) < a'_{k,r} \implies d_{k,r} < a'_{k,r} \leq t'_{k,r} \implies d_{k,r} < a_{k,r} \leq t_{k,r} \implies \min(d_{k,r}, t_{k,r}) < a_{k,r},$$

so the constraints bind in the truthful instance as well. Therefore, $a_{k,r}/e_k \leq x < x'$ for all agents k for which the first constraint binds in the misreported instance. Therefore x' is not the optimal objective value of the PSWC program in the misreported instance, a contradiction. Thus, $x \geq x'$. \square

Using Lemma 43, we show our main lemma. It allows us to make an inductive argument that, after gaining some extra resources in round r' , i 's allocation is (weakly) smaller for all other rounds in the misreported instance than the truthful instance.

LEMMA 44. *Let $r' < r < r'_i$ (that is, $a'_{i,r} < t'_{i,r}$). Suppose that $t_{j,r} \leq t'_{j,r}$ for all agents $j \neq i$. Then for all $j \neq i$, either: (1) $a_{j,r} = t_{j,r}$, or (2) $a_{j,r} \geq a'_{j,r}$.*

PROOF. Note that $t_{j,r} \leq t'_{j,r}$ for all $j \neq i$ implies that $t_{i,r} \geq t'_{i,r}$, which we use in the proof. Also, because $r' < r$, we know that $d'_{i,r} = d_{i,r}$, as r' is the last round for which $d'_{i,r} \neq d_{i,r}$. We assume that condition 1) from the lemma statement is false (i.e. $a_{j,r} < t_{j,r}$) and show that condition 2) must hold. Suppose first that $D'_r < E$. Then, because $a'_{i,r} < t'_{i,r}$, we know that $d'_{i,r} \leq t'_{i,r} \leq t_{i,r}$. This implies that $\min(d_{i,r}, t'_{i,r}) = \min(d_{i,r}, t_{i,r}) = d_{i,r}$. Let $j \neq i$. Since $t_{j,r} \leq t'_{j,r}$, we have $\min(d_{j,r}, t_{j,r}) \leq \min(d_{j,r}, t'_{j,r})$. Therefore, it is the case that $D_r \leq D'_r < E$. By Lemma 13 and the assumption that $a_{j,r} < t_{j,r}$, it must be the case that $a_{j,r} = \max(d_{j,r}, xe_j)$. Further, by Lemma 43, we know that $x \geq x'$. Therefore, we have

$$a'_{j,r} = \max(d_{j,r}, x'e_j) \leq \max(d_{j,r}, xe_j) = a_{j,r}.$$

That is, condition (2) from the lemma statement holds.

Now suppose that $D'_r \geq E$. Then, from the definition of the mechanism, we have that $a'_{j,r} \leq \min(d_{j,r}, t'_{j,r}) \leq d_{j,r}$. If it is the case that $D_r < E$ then we have that $a_{j,r} \geq \min(d_{j,r}, t_{j,r}) = d_{j,r}$, where the equality holds because otherwise we would have $a_{j,r} \geq \min(d_{j,r}, t_{j,r}) = t_{j,r}$, violating the assumption that $a_{j,r} < t_{j,r}$. Using these inequalities, we have $a_{j,r} \geq d_{j,r} \geq a'_{j,r}$, so condition (2) from the statement of the lemma holds. Finally, it may be the case that $D'_r \geq M$ and $D_r \geq M$. By Lemma 13 and the assumption that $a_{j,r} < t_{j,r}$, we have

$$a_{j,r} = \min(d_{j,r}, xe_k) \geq \min(d_{j,r}, x'e_k) = a'_{j,r},$$

where the inequality follows from Lemma 43. Thus, condition (2) of the lemma statement holds. \square

We now show an analogous result to Lemma 14.

LEMMA 45. *Suppose that $t_{j,r} \leq t'_{j,r}$ for all $j \neq i$, and $d_{k,r} = d'_{k,r}$ for all $k \in [n]$. Then $a_{i,r} \geq a'_{i,r}$.*

PROOF. Note that the condition that $t_{j,r} \leq t'_{j,r}$ for all $j \neq i$ implies that $t_{i,r} \geq t'_{i,r}$. We use these assumptions, along with the characterization of the FL mechanism allocations from Lemma 13, to prove the lemma.

We treat four cases, corresponding to whether or not supply exceeds demand in the truthful and misreported instances. Let x' denote the objective value in the FL mechanism's call to PSWC in the misreported instance, and x in the truthful instance. Suppose first that $D'_r \geq E$ and $D_r \geq E$. Suppose that $x \leq x'$. Then, for all $j \neq i$, $a_{j,r} = \min(xe_j, d_{j,r}, t_{j,r}) \leq \min(x'e_j, d_{j,r}, t'_{j,r}) = a'_{j,r}$, which implies that $a_{i,r} \geq a'_{i,r}$, since $\sum_{k \in [n]} a'_{k,r} = \sum_{k \in [n]} a_{k,r}$. On the other hand, if $x > x'$, then $a_{i,r} = \min(xe_i, d_{i,r}, t_{i,r}) \geq \min(x'e_i, d_{i,r}, t'_{i,r}) = a'_{i,r}$. Second, suppose that $D'_r \geq E$ and $D_r < E$. Then $a_{i,r} \geq \min(d_{i,r}, t_{i,r}) \geq \min(d_{i,r}, t'_{i,r}) \geq a'_{i,r}$. Third, suppose that $D'_r < E$ and $D_r \geq E$. Then $a_{j,r} \leq \min(d_{j,r}, t_{j,r}) \leq \min(d_{j,r}, t'_{j,r}) \leq a'_{j,r}$ for all $j \neq i$, which implies that $a_{i,r} \geq a'_{i,r}$.

Finally, suppose that $D'_r < E$ and $D_r < E$. If $x \leq x'$, then for all $j \neq i$, we have that $a_{j,r} = \min(t_{j,r}, \max(d_{j,r}, xe_j)) \leq \min(t'_{j,r}, \max(d_{j,r}, x'e_j)) = a'_{j,r}$, which implies that $a_{i,r} \geq a'_{i,r}$. If $x > x'$, then $a_{i,r} = \min(t_{i,r}, \max(d_{i,r}, xe_i)) \geq \min(t'_{i,r}, \max(d_{i,r}, x'e_i)) = a'_{i,r}$. Thus, the lemma holds in all cases. \square

Finally, we show that the mechanism is strategy-proof.

THEOREM 46. *Agent i never benefits from reporting $d_{i,r'} > d_{i,r}$.*

PROOF. We first observe that for every $r \leq r'_i$, $t_{j,r} \leq t'_{j,r}$ for every $j \neq i$. This is true for every $r \leq r'$ because $a'_{j,r} = a_{j,r}$ for $r < r'$, by Lemma 15. For $r = r' + 1$, it follows from Lemma 40, which says that $a_{j,r} \geq a'_{j,r'}$. For all subsequent rounds, up to and including $r = r'_i$, it follows inductively from Lemma 44: $t_{j,r} \leq t'_{j,r}$ implies that either $a_{j,r} = t_{j,r}$ (in which case $t_{j,r+1} = 0 \leq t'_{j,r+1}$), or $a_{j,r} \geq a'_{j,r}$ (in which case $t_{j,r+1} = t_{j,r} - a_{j,r} \leq t'_{j,r} - a'_{j,r} = t'_{j,r+1}$).

Consider an arbitrary round $r \neq r'$, with $r \leq r'_i$. By the above argument, we know that $t_{j,r} \leq t'_{j,r}$ for all $j \neq i$. Further, because reports in the truthful and misreported instances are identical on all rounds $r \neq r'$, we have that $d_{k,r} = d'_{k,r}$ for all $k \in [n]$. Therefore, by Lemma 45, $a_{i,r} \geq a'_{i,r}$. For rounds $r > r'_i$, it is also true that $a_{i,r} \geq a'_{i,r}$, since $a'_{i,r} = 0$ for these rounds by the definition of r'_i .

Finally,

$$\begin{aligned}
 U_{i,R}(\mathbf{a}_i) - U_{i,R}(\mathbf{a}'_i) &= \sum_{r=1}^R (u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r})) \\
 &= \sum_{r \neq r'} (u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r})) + (u_{i,r'}(a_{i,r'}) - u_{i,r'}(a'_{i,r'})) \\
 &= \sum_{r \neq r'} (u_{i,r}(a_{i,r}) - u_{i,r}(a'_{i,r})) - L(a'_{i,r'} - a_{i,r'}) \\
 &\geq L(a'_{i,r'} - a_{i,r'}) - L(a'_{i,r'} - a_{i,r'}) = 0
 \end{aligned}$$

Where the third transition follows from Corollary 42, and the final transition because $\sum_{r \neq r'} (a'_{i,r} - a_{i,r}) = a_{i,r'} - a'_{i,r'}$, and every term in the sum is positive. \square