Online EFX Allocations with Predictions

Extended Abstract

Themistoklis Melissourgos
University of Essex
Colchester, United Kingdom
themistoklis.melissourgos@essex.ac.uk

Nicos Protopapas
Archimedes Unit, Athena Research Center
Athens, Greece
n.protopapas@athenarc.gr

ABSTRACT

We study an online fair division problem where a fixed, but unknown number of goods arrive sequentially and must be allocated immediately and irrevocably to a given set of agents. The objective is to ensure (approximate) envy-freeness up to any good (EFX), that is, after the allocation, no agent should prefer another agent's bundle once any single good is removed from it. Unfortunately, we show that approximate EFX is impossible to guarantee, even under restrictive valuation assumptions.

To overcome this barrier, we follow the emerging trend of algorithms with predictions, assuming access to a vector of predicted valuations. Predictions may be inaccurate, and we measure their error using the total variation distance from the true valuations. For additive valuations, we prove impossibility results for algorithms that either ignore predictions or rely solely on them, and we establish lower bounds on the prediction accuracy required by any algorithm to compute approximate EFX. Finally, we provide a positive result: for two agents with identical valuations, we design an algorithm that uses predictions to achieve approximate EFX, with guarantees improving smoothly in prediction accuracy.

KEYWORDS

Online Fair Division, Envy-Freeness, Algorithms with Predictions, EFX

ACM Reference Format:
Themistoklis Melissourgos and Nicos Protopapas. 2026. Online EFX Allocations with Predictions: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/PGMH4700

1 INTRODUCTION

Fair division studies how shared resources can be allocated fairly among agents. A central fairness notion is *envy-freeness (EF)*, requiring that no agent prefers another's allocation. For divisible resources, EF allocations are guaranteed to exist under broad conditions [25]. For indivisible goods, however, even approximate EF may fail to exist: a single item desired by two agents can be allocated to only one of them. This impossibility has motivated the study of fairness relaxations that retain the spirit of EF.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/PGMH4700

One of the most prominent relaxations is *envy-freeness up to any good (EFX)* [10, 16]. While EFX is conjectured to always exist under monotone valuations, existence remains open even for four agents with additive valuations[1] and is only known for special cases [11]. Towards a proof of EFX existence, an approximation version, a-EFX, was introduced in [23]. Formally, an *allocation* $A = (A_1, A_2, \dots, A_n)$ of a set M of goods to n agents is a partition of M, where *bundle* A_i is allocated to agent $i \in [n]$. We say that A is a-EFX for some $a \in [0, 1]$ if there are no agents i, j and good $g \in A_j$ such that $v_i(A_i) < a \cdot v_i(A_j \setminus \{g\})$, and notice that for $a = 1$ we retrieve the (exact) EFX definition. Constant-factor guarantees are known [3, 5].

Beyond the classical offline setting, many allocation problems are inherently *online*: goods arrive sequentially and must be allocated to the agents immediately and irrevocably. Fairness is then evaluated only after all goods have been allocated. When future arrivals and values are unknown, fair allocation becomes significantly harder. In fact, recent work shows that in online settings without information about future goods, even guarantees weaker than EFX are impossible [22]. This motivates the question: *Can meaningful approximations of EFX be achieved online if limited, and possibly unreliable information about the future is available?*

We consider an online allocation model where agents provide predicted valuations for future goods, obtained for instance via historical data or machine-learning forecasts. Predictions may be inaccurate; as goods arrive, their true valuations are revealed, and the algorithm must allocate each good immediately. Our objective is to design online algorithms that leverage predictions and revealed values to guarantee a-EFX allocations, with a depending explicitly on the prediction accuracy.

Model and Notation. We study this question through the lens of *algorithmic design with predictions* [20]. Our setting is online, and involves discrete time-steps with a finite (true) horizon $T \geq 1$. Each agent i first receives a *prediction*, that is, a vector $p_i = (p_i(g_j))_{j \in [T']}$, where $T' \geq 1$ is the predicted horizon, and $p_i : M \to \mathbb{R}_{\geq 0}$ is an additive, normalized valuation function, i.e., $p_i(g_t) \geq 0$ for all $t \in [T']$, and $\sum_{t \in [T']} p_i(g_t) = 1$. Then, at each time $t = 1, 2, \dots, T$, a single good g_t arrives and has to be allocated *irrevocably* to some agent. Each agent $i \in [n]$ at time t evaluates good g_t according to his additive, normalized *true* valuation $v_i : M \to \mathbb{R}_{\geq 0}$, in other words, $v_i = (v_i(g_j))_{j \in [T]}$, where $v_i(g_t) \geq 0$ for all $t \in [T]$, and $\sum_{t \in [T]} v_i(g_t) = 1$. In the special case of *identical valuations* we drop the subscripts.

As a metric to quantify the prediction error, we use the *Total Variation distance (TV distance)* between the true and the predicted values: $\mathrm{TV}(p_i, v_i) = \|p_i - v_i\|_{\mathrm{TV}} := \frac{1}{2} \sum_{t \in [T_{\max}]} |p_i(g_t) - v_i(g_t)|$,

[1]Agent $i \in [n]$ has additive valuation $v_i : M \to \mathbb{R}_{\geq 0}$ if for any $S \subseteq M$, it holds that $v_i(S) = \sum_{g \in S} v_i(g)$.

where $T_{\max} := \max\{T, T'\}$.[2] We denote i's error by $d_i := \mathrm{TV}(p_i, v_i)$, and define $D := \max_{i \in [n]} d_i$. Alternatively, *accuracy* $\eta_i := 1 - d_i$, measures the fraction of the total value that is guaranteed to be predicted correctly by agent i.

1.1 Further Related Work

A body of related work studies alternative fairness notions [12, 24, 26, 27] with or without predictions, and considers settings where agents, rather than goods, arrive online [18, 19]; it also examines the online allocation of divisible resources [6, 7, 15]. A complementary line of work relaxes the information constraints of online models by allowing algorithms to peek into future arrivals [4, 17, 22], including the framework of *temporal fair division* [13, 14], where fairness must hold at every prefix—an especially restrictive requirement for EFX [14]. Overall, work on online fair division has studied similar core problems from different perspectives and under varying input assumptions [1, 2, 8, 9].

2 OUR RESULTS

We pose the question: *How does the quality (approximation factor a) of a-EFX allocations depend on the prediction accuracy of the agents?* We consider additive valuations and provide bounds on the prediction accuracy (or equivalently, on the allowed prediction error) as a function of $a \in [0, 1]$. Missing proofs and more results can be found in the full version of the paper [21].

We start by exploring the limitations of algorithms that do not have per-item predictions (but know the value of the whole set of goods for each agent, i.e., valuations are normalized). Our results show that for two agents with identical valuations, one can achieve $(\varphi - 1)$-EFX (where $\varphi := \frac{\sqrt{5}+1}{2}$) using a simple threshold-based algorithm, while no a-EFX algorithm without predictions exists for any $a \in (\varphi - 1, 1]$. When the valuations are not restricted to be identical, the latter impossibility result holds for any $a \in (0, 1]$.

Theorem 2.1. *Suppose we have 2 agents with additive, identical, normalized valuations, and without predictions. Consider the following algorithm: allocate each arriving good to agent 1, as long as the good will not make her exceed value $\varphi - 1$, otherwise allocate it to agent 2. This algorithm guarantees an allocation which is $(\varphi - 1)$-EFX. This bound is tight.*

Therefore, to achieve any improvement on a we turn our focus to algorithms with predictions. We seek algorithms that are consistent (i.e., guarantee high a when the predictions are perfect) and robust (i.e., guarantee high a when the predictions have unbounded error). Our result here leans on the negative: any algorithm that can guarantee $a > \varphi - 1 \approx 0.618$ with perfect predictions, cannot have robustness better than $(\varphi - 1)/2 \approx 0.309$. This means that, if our algorithms considered unbounded-error predictions, their guaranteed quality would be very low. Thus, we shift our focus to algorithms with *known bounds on the prediction error*.

We start by exploring the limitations of the extreme case, where the a-EFX algorithms rely entirely on predictions, disregarding the true values. The error (as a function of a) that these algorithms can withstand is a baseline for the error of general algorithms with predictions: the latter should tolerate at least that baseline error.

[2] If $T' < T$, then $p_i(g_t) = 0$ for $t \in \{T' + 1, T' + 2, \ldots, T\}$, while if $T' > 0$, then $v_i(g_t) = 0$ for $t \in \{T + 1, T + 2, \ldots, T'\}$.

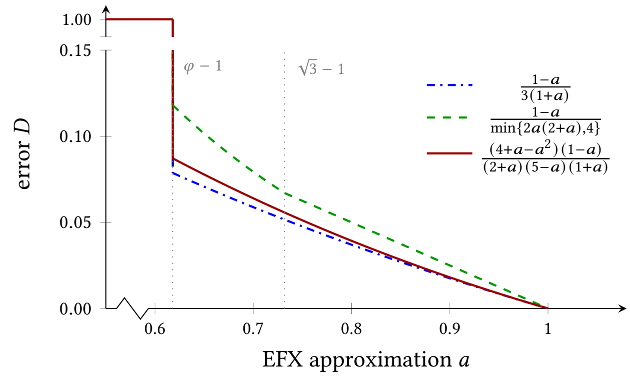

Figure 1: Maximum prediction error D as a function of $a \in [0, 1]$ for two agents with identical valuations. Dash-dot plot: Theorem 2.2, dashed plot: Theorem 2.3, solid plot: Theorem 2.1 and Theorem 2.4.

Theorem 2.2. *Suppose we have $n \geq 2$ agents with additive, normalized valuations. The agents have accuracy $\eta \geq 1 - \frac{\tilde{a}-a}{(2n-2+\tilde{a})(1+a)}$ for some given $a, \tilde{a} \in [0, 1]$ with $a \leq \tilde{a}$. Given an $\tilde{a}$-EFX allocation A^o according to $(p_i(g_t))_{i \in [n], t \in [T']}$, we can compute in polynomial time an allocation B which is a-EFX according to the true valuations $(v_i(g_t))_{i \in [n], t \in [T]}$. When $\tilde{a} = 1$, this accuracy is also necessary among algorithms oblivious to the true values.*

Next, we study the limitations of algorithms that use both predictions and true valuations, and show lower bounds on the level of accuracy as a function of the desired a. In particular, we show that for two agents, accuracy of $1 - \frac{1-a}{\min\{6a,4\}}$ is needed by any a-EFX algorithm for $a \in \left(\frac{1}{2}, 1\right]$. This bound slightly improves when the agents have identical valuations.

Theorem 2.3. *Suppose we have 2 agents with additive, identical, normalized valuations, with a provided prediction of accuracy $\eta < 1 - \frac{1-a}{\min\{2a(2+a),4\}}$ for some given $a \in (\varphi - 1, 1]$. Then, there is no algorithm that guarantees an a-EFX allocation, even when $T' = T = 4$.*

We show similarly strong bounds for $n \geq 3$ agents with identical valuations, even for $a \in (0, 1]$. Finally, we provide an algorithm in an attempt to bridge the gap between the lower bound $\left(1 - \frac{1-a}{\min\{2a(2+a),4\}}\right)$ and the upper bound $\left(1 - \frac{1-a}{3(1+a)}\right)$ on the accuracy for two agents with identical valuations.

Theorem 2.4. *Suppose we have 2 agents with additive, identical, normalized valuations, with a provided prediction of accuracy $\eta \geq 1 - \frac{(4+a-a^2)(1-a)}{(2+a)(5-a)(1+a)}$ for some given $a \in (\varphi - 1, 1]$. Then, there is an algorithm that outputs an a-EFX allocation, and performs a constant number of basic operations per time-step.*

Furthermore, this algorithm derives an improved upper bound of $1 - \frac{2}{5} \cdot \frac{1-a}{1+a}$ when predictions are 2-value functions, which we complement with a lower bound of $1 - \frac{1-a}{2}$. See Fig. 1 for a summary of our main results.

ACKNOWLEDGMENTS

This work has been partially supported by project MIS 5154714 of the National Recovery and Resilience Plan Greece 2.0 funded by the European Union under the NextGenerationEU Program.

REFERENCES

[1] Martin Aleksandrov, Haris Aziz, Serge Gaspers, and Toby Walsh. 2015. Online Fair Division: Analysing a Food Bank Problem. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI '15*. 2540–2546.

[2] Martin Aleksandrov and Toby Walsh. 2020. Online Fair Division: A Survey. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*. 13557–13562.

[3] Georgios Amanatidis, Aris Filos-Ratsikas, and Alkmini Sgouritsa. 2024. Pushing the Frontier on Approximate EFX Allocations. In *Proceedings of the 25th ACM Conference on Economics and Computation, EC '24*. 1268–1286.

[4] Georgios Amanatidis, Alexandros Lolos, Evangelos Markakis, and Victor Turmel. 2025. Online Fair Division for Personalized 2-Value Instances. In *Algorithmic Game Theory - 18th International Symposium, SAGT 2025, Proceedings*. 209–227.

[5] Georgios Amanatidis, Evangelos Markakis, and Apostolos Ntokos. 2020. Multiple birds with one stone: Beating 1/2 for EFX and GMMS via envy cycle elimination. *Theoret. Comput. Sci.* 841 (2020), 94–109.

[6] Siddhartha Banerjee, Vasilis Gkatzelis, Artur Gorokh, and Billy Jin. 2022. Online Nash social welfare maximization with predictions. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1–19.

[7] Siddhartha Banerjee, Vasilis Gkatzelis, Safwan Hossain, Billy Jin, Evi Micha, and Nisarg Shah. 2023. Proportionally Fair Online Allocation of Public Goods with Predictions. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI '23*. 20–28.

[8] Siddhartha Banerjee, Chamsi Hssaine, and Sean R. Sinclair. 2023. Online Fair Allocation of Perishable Resources. *SIGMETRICS Perform. Eval. Rev.* 51, 1 (2023), 55–56.

[9] Gerdus Benade, Aleksandr M Kazachkov, Ariel D Procaccia, Alexandros Psomas, and David Zeng. 2024. Fair and efficient online allocations. *Operations Research* 72, 4 (2024), 1438–1452.

[10] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Trans. Econ. Comput.* 7, 3 (2019).

[11] Bhaskar Ray Chaudhury, Jugal Garg, and Kurt Mehlhorn. 2024. EFX Exists for Three Agents. *J. ACM* 71, 1 (2024).

[12] Davin Choo, Winston Fu, Derek Khu, Tzeh Yuan Neoh, Tze-Yang Poon, and Nicholas Teh. 2025. Approximate Proportionality in Online Fair Division. arXiv:2508.03253

[13] Benjamin Cookson, Soroush Ebadian, and Nisarg Shah. 2025. Temporal fair division. *Proceedings of the AAAI Conference on Artificial Intelligence, AAAI '25* 39, 13 (2025), 13727–13734.

[14] Edith Elkind, Alexander Lam, Mohamad Latifian, Tzeh Yuan Neoh, and Nicholas Teh. 2025. Temporal Fair Division of Indivisible Items. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '25*. 676–685.

[15] Vasilis Gkatzelis, Alexandros Psomas, and Xizhi Tan. 2021. Fair and Efficient Online Allocations with Normalized Valuations. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 6 (2021), 5440–5447.

[16] Laurent Gourvès, Jérôme Monnot, and Lydia Tlilane. 2014. Near Fairness in Matroids. In *Proceedings of the Twenty-First European Conference on Artificial Intelligence, ECAI '14*. 393–398.

[17] Jiafan He, Ariel D. Procaccia, Alexandros Psomas, and David Zeng. 2019. Achieving a Fairer Future by Changing the Past. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI '19*. 343–349.

[18] Ian Kash, Ariel D Procaccia, and Nisarg Shah. 2014. No agent left behind: Dynamic fair division of multiple resources. *J. Artif. Intell. Res.* 51 (2014), 579–603.

[19] Pooja Kulkarni, Ruta Mehta, and Parnian Shahkar. 2025. Online Fair Division: Towards Ex-Post Constant MMS Guarantees. In *Proceedings of the 26th ACM Conference on Economics and Computation, EC '25*. 638–638.

[20] Thodoris Lykouris and Sergei Vassilvitskii. 2021. Competitive Caching with Machine Learned Advice. *J. ACM* 68, 4 (2021), 24:1–24:25.

[21] Themistoklis Melissourgos and Nicos Protopapas. 2025. Online EFX Allocations with Predictions. arXiv:2508.04779

[22] Tzeh Yuan Neoh, Jannik Peters, and Nicholas Teh. 2025. Online Fair Division with Additional Information. arXiv:2505.24503

[23] Benjamin Plaut and Tim Roughgarden. 2020. Almost envy-freeness with general valuations. *SIAM J. Discrete Math.* 34, 2 (2020), 1039–1068.

[24] Jiaxin Song, Biaoshuai Tao, Wenqian Wang, and Yuhao Zhang. 2025. Online MMS Allocation for Chores. arXiv:2507.14039

[25] Walter Stromquist. 1980. How to Cut a Cake Fairly. *Amer. Math. Monthly* 87, 8 (1980), 640–644.

[26] Yuanyuan Wang and Tianze Wei. 2026. Online Fair Allocations with Binary Valuations and Beyond. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI 2026)*. to appear.

[27] Shengwei Zhou, Rufan Bai, and Xiaowei Wu. 2023. Multi-agent online scheduling: MMS allocations for indivisible items. In *International Conference on Machine Learning*. 42506–42516.