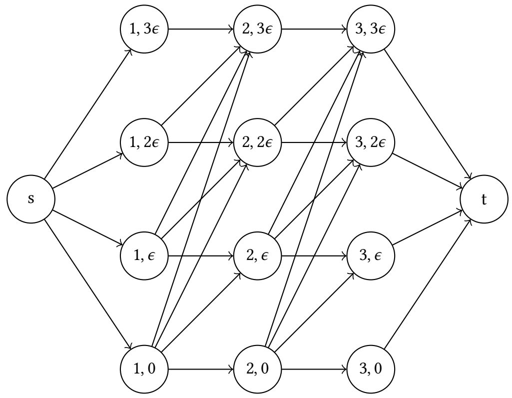

Figure 4: Example flow graph for $\epsilon = 1/3$ and three agents. A vertex v labeled "$i, y\epsilon$" indicates that y pieces of the resource have been consumed along the path from s to v.

We present the linear program for a single item below. Let $\mathcal{P} = \{2, 3, \ldots, n\} \times \{0, 1, 2, \ldots, 1/\epsilon\}$. The objective is again to maximize the total value of the agents. The first set of constraints impose flow preservation and the second constraint requires that a unit of flow is pushed from the source. The third set of constraints encodes ex-ante envy-freeness as agent i must weakly prefer her own lottery to that of i':

$$\text{maximize} \sum_{i=1}^{n} \sum_{j=0}^{1/\epsilon} \sum_{y=j}^{1/\epsilon} p_{i,j,y} f_i((y-j)\cdot\epsilon)$$

$$\text{subject to} \sum_{k=0}^{1/\epsilon} p_{i-1,k,j-k} = \sum_{\ell=0}^{1/\epsilon} p_{i,j,\ell}, \forall (i,j) \in \mathcal{P}$$

$$\sum_{j=0}^{1/\epsilon} p_{0,0,j} = 1$$

$$\sum_{j=0}^{1/\epsilon} \sum_{y=j}^{1/\epsilon} p_{i,j,y} f_i((y-j)\epsilon) \geq \sum_{j=0}^{1/\epsilon} \sum_{y=j}^{1/\epsilon} p_{i',j,y} f_i((y-j)\epsilon), \quad \forall i, i' \in \mathcal{N}$$

Theorem 5.1. *The linear program above outputs an ex-ante ϵ-Pareto optimal EF lottery in polynomial time using $O\left(\frac{mn}{\epsilon^2}\right)$ queries for any $\epsilon < \frac{1}{mn}$.*

Proof. First note that ex-ante envy-freeness is implied directly from the constraints of the program and that the lottery implied by the variables is a randomization over ex-post feasible outcomes by the discussion above. Let C denote the maximum Lipschitz constant across all pairs of agents and resources. To demonstrate that the lottery is ex-ante ϵ-Pareto optimal using at most $\frac{Cmn}{\epsilon^2}$ queries, first observe that the lottery which gives each agent all the items with probability $1/n$ is a feasible solution to our program and is consistent with our discretization. On the other hand, as the objective function is optimizing the sum of ex-ante utilities (subject to the constraints), the linear program outputs a lottery which is ex-ante Pareto optimal with respect to the discretization. Therefore, the linear program must give at least one agent utility greater than or equal to $1/n$. The only reason that the solution for our program would not be ex-ante Pareto efficient with respect to the full valuation functions would be if there were jumps in the valuation functions between our queries. However, since all agents have Lipschitz valuation functions with Lipschitz constant at most C, in any ex-post outcome, if we were to give every agent an additional piece of size $\frac{\epsilon^2}{C}$ of every item, the utility of an agent would increase by at most $m\epsilon^2$. For an agent with utility $v \geq 1/n$, we have $v \cdot (1+\epsilon) \geq v + \frac{\epsilon}{n} > v + m\epsilon^2$. So, there is no lottery that improves her utility by a factor $1+\epsilon$, even with respect to the exact valuation functions. Finally, since there are a polynomial number of variables and constraints the program runs in polynomial time. □

We also note that our algorithm for finding ex-ante approximate Pareto efficient EF lotteries can be extended to agents with *non-additive*, bounded gradient valuation functions for a constant number of items. The procedure is largely the same, except we ask $\frac{1}{\epsilon^m}$ value queries for all possible combinations of ϵ sized pieces, e.g., for two items $(0,0), (0,\epsilon), (0,2\epsilon), \ldots, (0,1), (\epsilon,0), (\epsilon,\epsilon), \ldots, (1,1)$. We then appropriately modify the flow graph to include vertices representing all possible tuples for each agent and direct edges accordingly. However, this procedure only works for a constant number of items as the number of queries then depends *exponentially* on the number of items.

Since we discretize the resources into $1/\epsilon$ pieces each and treat these pieces as indivisible items, one could consider using the approach of Budish et al. [11] to find a lottery over ex-post feasible outcomes. However, our valuation functions over the items can be more general than their approach can support since we allow complementarity within a single resource (e.g., by allowing for convex valuation functions). On the other hand, their approach allows for more expressive constraint structures over the feasible allocations of goods to agents. We believe that investigating the degree to which the two approaches can be combined is an interesting question for future work.

6 OBSERVATIONS AND FUTURE WORK

In this section, we provide some auxiliary results in our model and suggest several interesting related open questions. An interesting first question to consider would be to tighten the gap between the upper and lower bounds on the number of queries required to compute an ex-ante ϵ-Pareto efficient EF lottery for a constant number of players. Do $o\left(\frac{1}{\epsilon^2}\right)$ queries suffice? In the case of general m and n, do $o\left(\frac{mn}{\epsilon^2}\right)$ suffice?

While our lower bound suggests that exact ex-ante Pareto efficient lotteries are unattainable using a bounded number of queries, if we simplify our objective from ex-ante Pareto efficiency to ex-post Pareto efficiency then we may find an ex-ante envy-free lottery which is *exactly* ex-post Pareto efficient with respect to the set of all outcomes. We note that our algorithm closely resembles the well-known random serial dictatorship mechanism for one-sided matching markets of Abdulkadiroğlu and Sönmez [1].

Theorem 6.1. *There exists an algorithm that produces an ex-ante envy-free lottery that is exactly ex-post Pareto efficient among the set of all outcomes using a polynomial number of queries.*

Proof. Consider an arbitrary set of n agents $\mathcal{N}$. We uniformly and randomly permute the agents and label them $1, 2, \ldots, n$. We ask agent 1 a value query $\text{Value}(f_{1k}, 1)$ for each item $k \in \mathcal{M}$. In doing so, we determine her value v_{1k} for receiving each good k entirely. We then ask her a cut query $\textsc{Cut}(f_{1k}, v_{1k})$ for each item $k \in \mathcal{M}$, thereby determining how much of each good she needs to receive to obtain her "full value". We continue by allocating to her as much of each item as the corresponding cut query returns. For any remaining pieces of each item, we repeat this process for each successive agent in the ordering. As valuation functions are additive across items, it is easy to verify that this is ex-post Pareto efficient; removing any amount of any of the goods from agent i must decrease her value. This lottery is also ex-ante envy-free, since every agent has the same probability of appearing at any index in the random ordering, and at every point in the algorithm each consecutive agent receives a favorite bundle of goods among the remaining resources. Therefore, no agent envies the lottery of any other agent. □

On the other hand, while ex-post Pareto efficiency is easier to achieve than ex-ante Pareto efficiency, ex-post envy-freeness is a more stringent requirement than ex-ante envy-freeness. For a lottery to be ex-post envy-free, it must be that all outcomes in its support are ex-post envy-free. One may then ask the same questions we explore in this paper through the context of the class of ex-post envy-free lotteries. For instance, can one obtain an ex-ante ϵ-Pareto efficient lottery among the set of all ex-post envy-free lotteries using only a polynomial number of queries?

One can also imagine asking a similar question of finding ex-ante Pareto optimal lotteries among the set of all *ex-ante proportional* lotteries. Ex-ante proportionality is another standard notion of fairness which requires all agents to receive expected utility at least $1/n$. By changing the third set of constraints in our linear program in Section 5 to encode proportionality, one directly obtains the same guarantee of Theorem 5.1, i.e., one obtains an approximately ex-ante Pareto optimal proportional lottery. However, perhaps even more excitingly, one could change the welfare-maximization objective and instead maximize the *leximin* utility by repeatedly maximizing the minimum value (as in, e.g., Kurokawa et al. [20]) and *drop the fairness constraint altogether*. Since the equiprobable lottery is in the space of feasible lotteries, the leximin solution will necessarily give all agents expected utility at least $1/n$. Thus, the lottery would be approximately ex-ante Pareto optimal among *all* lotteries and satisfy ex-ante proportionality!

We view the question of determining if one can find a lottery which is approximately ex-ante Pareto optimal among all lotteries and which is also ex-ante envy-free to be a very interesting and challenging question. In fact, Brams et al. [8] point to the related question of finding a tractable algorithm to compute welfare-maximizing envy-free and Pareto efficient allocations (among the set of all allocations) in the standard cake-cutting model even in the seemingly quite simple case of known piecewise constant valuations as their "most important, and presumably quite challenging, open problem". More broadly, investigating the extent to which different fairness notions change the shape of the ex-ante Pareto frontier is an interesting line of future study.

7 CONCLUSION

In this paper, we propose a new model for allocating homogeneous divisible goods and we show that it is simple in natural ways – envy-freeness is easy to achieve – yet highly complex in others – many descriptive, non-trivial value functions can be captured. We provide an algorithm for finding approximately ex-ante Pareto optimal EF lotteries and complement this result with a lower bound on the query complexity of doing so. Notably, while our algorithm uses only value queries, our lower bound applies to any algorithm using either cut and/or value queries. We also provide an algorithm for finding exactly ex-post Pareto optimal EF lotteries which uses both cut and value queries.

Throughout this paper, we assume that valuation functions are Lipschitz. Therefore, determining if approximate Pareto optimality can be achieved when the Lipschitz assumption is removed would be a first step in extending our results. Perhaps most importantly, removing the assumption that agents are additive across resources is a major open line of further work. We have a way of converting our approach for an arbitrary number of additive items to a computationally efficient method of finding approximately ex-ante Pareto efficient EF lotteries for general valuations if there is a constant number of items. On the other hand, extending this result to arbitrarily many items may require very different techniques.

ACKNOWLEDGMENTS

The second author was supported in part by an NSF CAREER award CCF-2047907 and NSF grant CCF-2008280. The third author was supported in part by a Google Research Scholar Award. The fourth author was supported in part by NSF grants CCF-2008280 and CCF-1755955. Part of this work was done while the last three authors were visiting the Simons Institute for the Theory of Computing.

REFERENCES

[1] Atila Abdulkadiroğlu and Tayfun Sönmez. 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66, 3 (1998), 689–701.

[2] Georgios Amanatidis, George Christodoulou, John Fearnley, Evangelos Markakis, Christos-Alexandros Psomas, and Eftychia Vakaliou. 2018. An improved envy-free cake cutting protocol for four agents. In *Proceedings of the 11th International Symposium on Algorithmic Game Theory (SAGT) (Lecture Notes in Computer Science, Vol. 11059).* 87–99.

[3] Haris Aziz and Simon Mackenzie. 2016. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS).* 416–427.

[4] Haris Aziz and Simon Mackenzie. 2016. A discrete and bounded envy-free cake cutting protocol for four agents. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC).* 454–464.

[5] Xiaohui Bei, Ning Chen, Xia Hua, Biaoshuai Tao, and Endong Yang. 2012. Optimal proportional cake cutting with connected pieces. In *Twenty-sixth AAAI Conference on Artificial Intelligence.*

[6] Garrett Birkhoff. 1946. Three observations on linear algebra. *Univ. Nac. Tacuman, Rev. Ser. A* 5 (1946), 147–151.

[7] Anna Bogomolnaia and Hervé Moulin. 2001. A new solution to the random assignment problem. *Journal of Economic Theory* 100, 2 (2001), 295–328.

[8] Steven Brams, Michal Feldman, John Lai, Jamie Morgenstern, and Ariel Procaccia. 2012. On maxsum fair cake divisions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 26.

[9] S. Brams and A. Taylor. 1996. *Fair Division - From Cake-Cutting to Dispute Resolution.* Cambridge University Press.

[10] Steven J Brams and Alan D Taylor. 1995. An envy-free cake division protocol. *The American Mathematical Monthly* 102, 1 (1995), 9–18.

[11] Eric Budish, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom. 2013. Designing random allocation mechanisms: Theory and applications. *American Economic Review* 103, 2 (2013), 585–623.

[12] Jan Buermann, Enrico H Gerding, and Baharak Rastegari. 2020. Fair allocation of resources with uncertain availability. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS).* 204–212.

[13] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. 2012. The efficiency of fair division. *Theory of Computing Systems*

50, 4 (2012), 589–610.
[14] Yuga J Cohler, John K Lai, David C Parkes, and Ariel D Procaccia. 2011. Optimal envy-free cake cutting. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI).* 626–631.
[15] Lester E Dubins and Edwin H Spanier. 1961. How to cut a cake fairly. *The American Mathematical Monthly* 68, 1P1 (1961), 1–17.
[16] Uriel Feige and Moshe Tennenholtz. 2014. On fair division of a homogeneous good. *Games and Economic Behavior* 87 (2014), 305–321.
[17] Eric J Friedman, Vasilis Gkatzelis, Christos-Alexandros Psomas, and Scott Shenker. 2019. Fair and efficient memory sharing: Confronting free riders. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 1965–1972.
[18] Nicole Immorlica, Brendan Lucier, E. Glen Weyl, and Joshua Mollner. 2017. Approximate Efficiency in Matching Markets. In *Web and Internet Economics - 13th International Conference, WINE 2017, Bangalore, India, December 17-20, 2017, Proceedings (Lecture Notes in Computer Science, Vol. 10660)*, Nikhil R. Devanur and Pinyan Lu (Eds.). Springer, 252–265.
[19] David Kurokawa, John K. Lai, and Ariel D. Procaccia. 2013. How to Cut a Cake before the Party Ends. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)* (Bellevue, Washington) *(AAAI'13)*. AAAI Press, 555–561.
[20] David Kurokawa, Ariel D. Procaccia, and Nisarg Shah. 2018. Leximin Allocations in the Real World. *ACM Trans. Econ. Comput.* 6, 3–4, Article 11 (Oct. 2018), 24 pages. https://doi.org/10.1145/3274641
[21] Ariel D Procaccia. 2009. Thou shalt covet thy neighbor's cake. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI).*
[22] J Hans Reijnierse and Jos AM Potters. 1998. On finding an envy-free Pareto-optimal division. *Mathematical Programming* 83, 1-3 (1998), 291–311.
[23] J. M. Robertson and W. Webb. 1998. *Cake-Cutting Algorithms - Be Fair If You Can.* CRC Press.
[24] Günther Ruhe and Bernd Fruhwirth. 1990. ε-optimality for bicriteria programs and its application to minimum cost flows. *Computing* 44, 1 (1990), 21–34.
[25] John Von Neumann. 2016. 1. A Certain Zero-sum Two-person Game Equivalent to the Optimal Assignment Problem. In *Contributions to the Theory of Games (AM-28), Volume II*. Princeton University Press, 5–12.
[26] David Zeng and Alexandros Psomas. 2020. Fairness-Efficiency Tradeoffs in Dynamic Fair Division. In *EC '20: The 21st ACM Conference on Economics and Computation, Virtual Event, Hungary, July 13-17, 2020*, Péter Biró, Jason Hartline, Michael Ostrovsky, and Ariel D. Procaccia (Eds.). ACM, 911–912.