The allocation A is EF if and only if E_{ij} does not occur for all $i \neq j$. Using Equation (3) and the union bound over $\binom{n}{2}$ pairs of agents, the probability that A is *not* EF is at most

$$\Pr\left[\bigvee_{i \neq j} E_{ij}\right] \leq \sum_{i \neq j} \Pr[E_{ij}] \leq \binom{n}{2} \frac{1}{m^3} \leq \frac{1}{m}.$$

Thus, the probability that A is not EF goes to zero as m grows. □

In Between: A Phase Transition

In this section, we support our theoretical results with an empirical exploration of the transition from nonexistence to existence of envy-free allocations as a function of the number of goods and agents. We find that the most difficult allocation problems occur during the sharp *phase transition* from nonexistence to existence. We show that this behavior, which is common to many discrete feasibility problems, holds under both of two natural optimization models (one with and one without an objective function) and under different distributions over agents' utility values.

Experimental Setup

We generate instances with n agents and m goods as follows by sampling valuations for each agent and each good from a given distribution over utility functions. In our experimental setup, we draw from two distributions—CORRELATED$(0.4, 0.6)$ and UNIFORM$(0, 1)$—defined earlier. Intuitively, the UNIFORM distribution randomly assigns a value to each good for each agent, while the CORRELATED distribution first draws an intrinsic value for each good, then assigns a random value to each agent drawn from a (truncated nonnegative normal) distribution around that intrinsic value. UNIFORM satisfies both distributional assumptions and thus aligns with both Theorems 1 and 2, while our instantiation of CORRELATED only satisfies assumption [A2], or the assumption needed for Theorem 2. Still, we will show that both theoretical results hold experimentally for both distributions, even when the number of agents and goods is quite small.

Given an instance as generated above, we search for an envy-free allocation using one of two mixed integer programs (MIPs). Both formulations use $n \times m$ binary variables x_{ig} that are activated if and only if agent i is allocated good g. Model #1, a feasibility problem, is defined as follows:

$$\begin{array}{lll} \text{find} & x_{ig} \quad \forall i \in N, g \in G & \\ \text{s.t.} & \sum_{i \in N} x_{ig} = 1 & \forall g \in G \\ & \sum_{g \in G} v_{ig} x_{i'g} - \sum_{g \in G} v_{ig} x_{ig} \leq 0 & \forall i \neq i' \in N \\ & x_{ig} \in \{0, 1\} & \forall i \in N, g \in G \end{array}$$

Intuitively, the first set of constraints ensures that each good is allocated to exactly one agent, while the second set of constraints ensures that each agent values its allocation at least as highly as any other agent's allocation. For this feasibility problem, no explicit objective function is necessary; indeed, the feasible region defined by the constraints is exactly the space of all envy-free allocations.

We now define Model #2, an optimization version of the envy-free allocation problem, as follows:

$$\begin{array}{lll} \min & e & \\ \text{s.t.} & \sum_{i \in N} x_{ig} = 1 & \forall g \in G \\ & \sum_{g \in G} v_{ig} x_{i'g} - \sum_{g \in G} v_{ig} x_{ig} \leq e & \forall i \neq i' \in N \\ & x_{ig} \in \{0, 1\} & \forall i \in N, g \in G \\ & e \in \mathbb{R}^{\text{nonneg}} & \end{array}$$

This second MIP model minimizes a real-valued nonnegative variable e representing the maximum envy between any two agents; thus, an EF allocation exists if and only if the objective value is zero at the optimum. This is an integer programming-based implementation of the envy minimization problem described by Lipton et al. (2004).

Model #1 may seem like the more general model since it is amenable to the addition of various objective functions. For example, adding an objective function that maximizes $\sum_{i \in N} \sum_{g \in G} v_{ig} x_{ig}$ would produce an envy-free allocation that also maximizes social welfare. It is not obvious how to adapt Model #2 to include arbitrary objective functions. Still, there is some evidence that relaxing the feasible region and then re-casting the feasibility problem as an optimization problem may result in better runtime performance. For example, Sandholm, Gilpin, and Conitzer (2005) saw speedups using an optimization model instead of a feasibility model in specific problem classes when exploring various MIP models for finding Nash equilibria in two-player games (although they did not see an overall speedup). We compare the performance of both models in the coming section.

All experiments were performed in Python using IBM ILOG CPLEX 12.6[1] in single-threaded mode under its default configuration.[2] Runs were conducted on Blacklight,[3] a ccNUMA supercomputer with 8GB of RAM per core; each experiment was run at least 160 times with a time limit of 12 hours per run. For solve time comparison, runs that timed out were conservatively considered to have completed in 12 hours. When timeouts were ignored or penalized heavily (e.g., counted as a $10 \times 12 = 120$ hour run), our experiments exhibited the same qualitative behavior.

Phase Transitions

We now explore the existence of phase transitions in various instantiations of the envy-free allocation problem.

Figure 1 shows an example phase transition for the existence of, and hardness of finding, an envy-free allocation in a problem with $n = 10$ agents valuing $m \in \{10, \ldots, 30\}$ goods. Results are presented for both the UNIFORM and CORRELATED distributions over utility functions using Model #1 without and with a social welfare maximizing objective function. The thick red line (corresponding to the left y-axis) plots the fraction of instances with m goods and n agents such that an envy-free allocation existed.

Aligning with Theorem 1, Figure 1 shows that the probability of an EF allocation existing is small when the number of goods is not much larger than the number of agents. Similarly, aligning with Theorem 2, when the number of goods is

[1] ibm.com/software/commerce/optimization/cplex-optimizer/

[2] Source code & data: https://github.com/JohnDickerson/EnvyFree

[3] blacklight.psc.edu

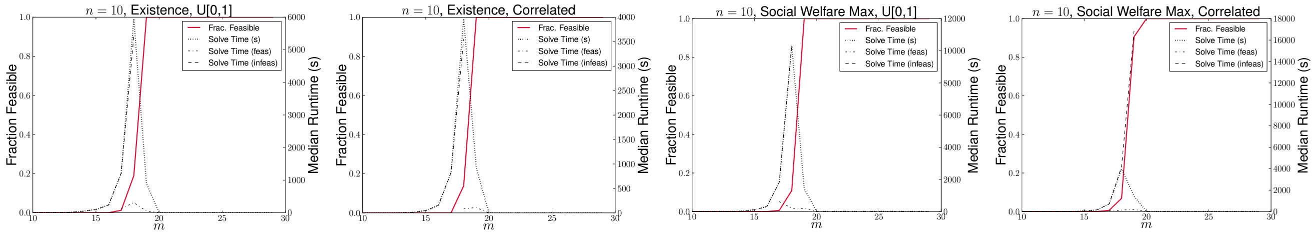

Figure 1: Phase transition for $n = 10$ under either UNIFORM or CORRELATED, with or without maximizing social welfare.

more (but not necessarily substantially more), the probability of an EF allocation existing is essentially one. Figure 2 explores this transition quantitatively for increasing numbers of agents n by plotting the minimum value m where at least 99% of the generated instances were feasible. Fitting an $m/\ln(m)$ function for either UNIFORM or CORRELATED shows that the asymptotically-stated Theorem 2 holds even when the number of goods and agents is quite small.

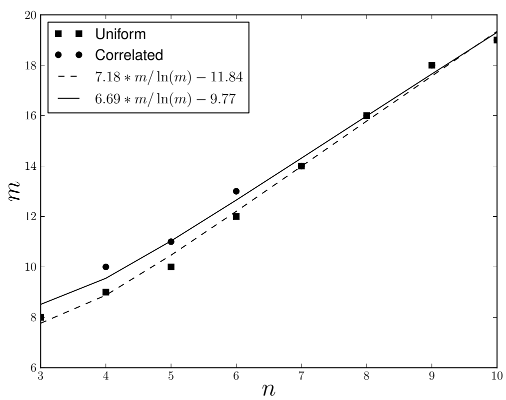

Figure 2: The minimum value of m where at least 99% of the instances were feasible as n increases.

Figure 1 also plots runtime as a function of the number of goods m. The thick dashed line (corresponding to the right y-axis) plots the median runtime to either prove the nonexistence of a solution or find and prove the optimality of a feasible solution. The two dotted lines (also corresponding to the right y-axis) plot the median runtimes for only the feasible and infeasible instances, respectively. We see a classical "hardness bump" around the phase transition, with median solution time being much higher when the probability of a feasible instance is small but not trivial. Here, proving infeasibility takes significant computational effort.

Figure 3 shows that this hardness behavior is not just an artifact of the feasibility Model #1; indeed, the optimization problem defined by Model #2 exhibits an even more stark hardness bump around the phase transition. This roughly aligns with the experiences of Sandholm, Gilpin, and Conitzer (2005), who found that relaxing the feasible region while moving some constraints into the objective did not result in an overall speedup.

Discussion & Future Research

In this paper, we theoretically and empirically investigated the existence of envy-free allocations of indivisible goods. Under additive valuations and general assumptions on the distributions over values of individual goods, we theoretically characterized the conditions for nonexistence and existence of envy-free allocations. We supported these asymptotic results with experiments on two value distributions using two MIP models and found, empirically, that the theoretical conditions for (non)existence of envy-free allocations apply even when the number of agents and goods is quite small. Furthermore, we discovered that the hardest computational problems in this space on average exist during the phase transition between nonexistence and existence.

In typical phase transition work, what is increased on the "x-axis" is the number of constraints while keeping the number of variables constant. Our phase transition is, in that sense, different because as we increase the number of goods (while keeping the number of agents fixed), both the number of variables and constraints increases. Our phase transition is nevertheless similar to prior ones in that (i) there is a sharp transition from infeasibility to feasibility, (ii) the complexity peak occurs at that transition, (iii) the complexity peak is driven mainly by infeasible instances, and (iv) the infeasible instances get harder—and rarer—as we move to the side of the phase transition where instances are typically feasible.

While the theoretical results we presented are essentially tight, it would be useful to completely characterize the phase transition between nonexistence and existence of an envy-free allocation. We showed experimentally that this phase transition is quite sharp, but either proving that the logarithmic factor in Theorem 2 is necessary or further whittling down this bound toward Theorem 1 would be helpful. Results of this nature are actively being pursued with random 3-SAT problems (Kaporis, Kirousis, and Lalas 2006; Maneva and Sinclair 2008). Furthermore, relaxing the distributional assumptions (especially on Theorem 1) would, if possible, be useful toward this end.

Along the lines of enhanced MIP techniques, it would be interesting to try to "flatten the hardness bump" we saw in the experiments through the use of custom branching and fathoming rules, variable prioritization schemes, and other heuristics that maintain search completeness.

Acknowledgments

This material was funded by NSF grants IIS-1320620, CCF-1101668, CCF-1215883, and IIS-0964579, by an NDSEG fellowship, and used the Pittsburgh Supercomputing Center in partnership with the XSEDE, which is supported by NSF grant OCI-1053575. We thank Intel Corporation for machine gifts.

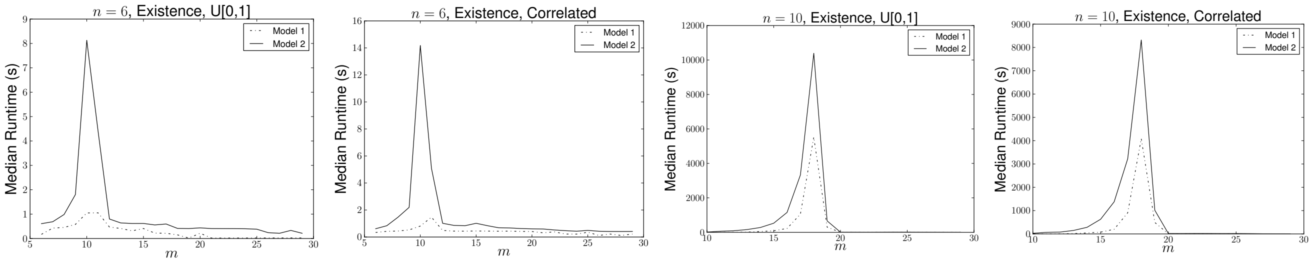

Figure 3: Runtime comparison of Model #1 (feasibility) and Model #2 (optimization) for $n = 6$ and $n = 10$ agents.

References

Boutilier, C.; Caragiannis, I.; Haber, S.; Lu, T.; Procaccia, A. D.; and Sheffet, O. 2012. Optimal social choice functions: A utilitarian view. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, 197–214.

Bouveret, S., and Lang, J. 2008. Efficiency and envy-freeness in fair division of indivisible goods: logical representation and complexity. *Journal of Artificial Intelligence Research* 32:525–564.

Brams, S. J., and Fishburn, P. C. 2000. Fair division of indivisible items between two people with identical preferences: Envy-freeness, Pareto-optimality, and equity. *Social Choice and Welfare* 17:247–267.

Brams, S. J.; Kilgour, M.; and Klamler, C. 2014. Two-person fair division of indivisible items: An efficient, envy-free algorithm. *Notices of the AMS* 61(2):130–141.

Cheeseman, P.; Kanefsky, B.; and Taylor, W. 1991. Where the really hard problems are. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI)*, 331–337.

Chen, Y.; Lai, J. K.; Parkes, D. C.; and Procaccia, A. D. 2010. Truth, justice, and cake cutting. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, 756–761.

Chevaleyre, Y.; Dunne, P. E.; Endriss, U.; Lang, J.; Lemaître, M.; Maudet, N.; Padget, J.; Phelps, S.; Rodríguez-Aguilar, J. A.; and Sousa, P. 2006. Issues in multiagent resource allocation. *Informatica* 30:3–31.

Chevaleyre, Y.; Endriss, U.; and Maudet, N. 2007. Allocating goods on a graph to eliminate envy. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI)*, 700–705.

Cohler, Y. J.; Lai, J. K.; Parkes, D. C.; and Procaccia, A. D. 2011. Optimal envy-free cake cutting. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI)*, 626–631.

Conitzer, V., and Sandholm, T. 2006. Nonexistence of voting rules that are usually hard to manipulate. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*.

Conitzer, V.; Sandholm, T.; and Lang, J. 2007. When are elections with few candidates hard to manipulate? *Journal of the ACM* 54(3):1–33.

Hogg, T.; Huberman, B. A.; and Williams, C. P. 1996. Phase transitions and the search problem. *Artificial Intelligence* 81(1–2):1–15.

Kaporis, A. C.; Kirousis, L. M.; and Lalas, E. G. 2006. The probabilistic analysis of a greedy satisfiability algorithm. *Random Structures & Algorithms* 28(4):444–480.

Lipton, R.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *ACM Conference on Electronic Commerce (EC)*, 125–131.

Maneva, E., and Sinclair, A. 2008. On the satisfiability threshold and clustering of solutions of random 3-sat formulas. *Theoretical Computer Science* 407(1):359–369.

Mossel, E.; Procaccia, A. D.; and Rácz, M. Z. 2013. A smooth transition from powerlessness to absolute power. *Journal of Artificial Intelligence Research* 48:923–951.

Procaccia, A. D., and Rosenschein, J. S. 2007. Average-case tractability of manipulation in elections via the fraction of manipulators. In *Proceedings of the 6th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 718–720.

Procaccia, A. D. 2013. Cake cutting: Not just child's play. *Communications of the ACM* 56(7):78–87.

Sandholm, T.; Gilpin, A.; and Conitzer, V. 2005. Mixed-integer programming methods for finding Nash equilibria. In *Proceedings of the 20th AAAI Conference on Artificial Intelligence (AAAI)*, 495–501.

Walsh, T. 2011. Where are the hard manipulation problems? *Journal of Artificial Intelligence Research* 42:1–29.

Xia, L., and Conitzer, V. 2008. Generalized scoring rules and the frequency of coalitional manipulability. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*, 109–118.