

B. REFERENCES

- [1] Mechanical turk platform. <https://www.mturk.com/>.
- [2] A. Archer and E. Tardos. Frugal path mechanisms. In *SODA*, pages 991–999, 2002.
- [3] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2-3):235–256, 2002.
- [4] P. Auer, R. Ortner, and C. Szepesvari. Improved rates for the stochastic continuum-armed bandit problem. In *COLT*, pages 454–468, 2007.
- [5] M. Babaioff, M. Dinitz, A. Gupta, N. Immorlica, and K. Talwar. Secretary problems: weights and discounts. In *SODA*, pages 1245–1254, 2009.
- [6] M. Babaioff, S. Dughmi, R. Kleinberg, and A. Slivkins. Dynamic pricing with limited supply. In *EC*, 2012.
- [7] M. Babaioff, N. Immorlica, D. Kempe, and R. Kleinberg. A knapsack secretary problem with applications. In *APPROX-RANDOM*, pages 16–28, 2007.
- [8] A. Badanidiyuru, R. Kleinberg, and Y. Singer. Learning on a budget: Posted price mechanisms for online procurement. In *EC*, 2012.
- [9] Z. Bar-yossef, K. Hildrum, and F. Wu. Incentive-compatible online auctions for digital goods. In *SODA*, pages 964–970, 2002.
- [10] O. Besbes and A. Zeevi. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research*, 2009.
- [11] A. Blum and J. Hartline. Near-optimal online auctions. In *SODA*, 2005.
- [12] A. Blum, V. Kumar, A. Rudra, and F. Wu. Online learning in online auctions. In *SODA*, 2003.
- [13] A. Blum and Y. Mansour. *Learning, Regret minimization, and Equilibria in Algorithmic Game Theory*. Cambridge University Press, 2007.
- [14] M. C. Cary, A. D. Flaxman, J. D. Hartline, and A. R. Karlin. Auctions for structured procurement. In *SODA*, pages 304–313, 2008.
- [15] N. Cesa-Bianchi, Y. Freund, D. P. Helmbold, D. Haussler, R. Schapire, and M. Warmuth. How to use expert advice. *Journal of the ACM*, 44(2):427–485, 1997.
- [16] S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Multi-parameter mechanism design and sequential pricing. In *EC*, 2010.
- [17] Y. Chen and J. W. Vaughan. A new understanding of prediction markets via no-regret learning. In *STOC*, 2010.
- [18] N. Devanur and J. Hartline. Limited and online supply and the bayesian foundations of prior-free mechanism design. In *EC*, 2009.
- [19] A. Goldberg, J. Hartline, and A. Wright. Competitive auctions and digital goods. In *SODA*, pages 735–744, 2001.
- [20] M. T. Hajiaghayi, R. D. Kleinberg, and D. C. Parkes. Adaptive limited-supply online auctions. In *EC*, 2004.
- [21] J. J. Horton and L. B. Chilton. The labor economics of paid crowdsourcing. In *EC*, pages 209–218, 2010.
- [22] J. J. Horton and R. J. Zeckhauser. Algorithmic wage negotiations: Applications to paid crowdsourcing. In *CrowdConf*, 2010.
- [23] R. Karlin, D. Kempe, and T. Tamir. Beyond vcg: Frugality of truthful mechanisms. In *FOCS*, 2005.
- [24] R. Kleinberg. Nearly tight bounds for the continuum-armed bandit problem. In *NIPS*, 2004.
- [25] R. D. Kleinberg and F. T. Leighton. The value of knowing a demand curve: bounds on regret for online posted-price auctions. In *FOCS*, 2003.
- [26] R. Lavi and N. Nisan. Competitive analysis of incentive compatible on-line auctions. In *EC*, pages 233–241, 2000.
- [27] N. Littlestone and M. K. Warmuth. The weighted majority algorithm. *Info and Computation*, 70(2):212–261, 1994.
- [28] W. Mason and D. J. Watts. Financial incentives and the performance of crowds. In *SIGKDD Explor. Newsl.*, pages 100–108, 2010.
- [29] J. J. Shaw, A. D.; Horton and D. L. Chen. Designing incentives for inexperienced human raters. In *CSCW*, pages 275–284, 2011.
- [30] Y. Singer. Budget feasible mechanisms. In *FOCS*, pages 765–774, 2010.
- [31] Y. Singer. How to win friends and influence people, truthfully: Influence maximization mechanisms for social networks. In *WSDM*, 2011.
- [32] Y. Singer and M. Mittal. Pricing tasks in online labor markets. In *Workshop on Human Computation*. ACM, 2011.
- [33] N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *ICML*, 2010.
- [34] Tran-Thanh, L., C. A., J. E. M. de Cote, A. Rogers, and N. R. Jennings. Epsilon-first policies for budget-limited multi-armed bandits. In *AAAI*, pages 1211–1216, 2010.
- [35] L. Tran-Thanh, A. Chapman, A. Rogers, and N. R. Jennings. Knapsack based optimal policies for budget-limited multi-armed bandits. <http://arxiv.org/abs/1204.1909>, 2013. Technical report.
- [36] L. Tran-Thanh, A. Chapman, A. Rogers, and N. R. Jennings. Knapsack based optimal policies for budget-limited multi-armed bandits. In *AAAI*, 2012.
- [37] L. Tran-Thanh, S. Stein, A. Rogers, and N. R. Jennings. Efficient crowdsourcing of unknown experts using multi-armed bandits. In *AAAI*, 2012.