Quick or Cheap? Breaking Points in Dynamic Markets

PANAYOTIS MERTIKOPOULOS, CNRS, France – Grenoble, France
HEINRICH H. NAX, University of Zurich, Switzerland
BARY S. R. PRADELSKI, CNRS, France - Grenoble, France

We examine two-sided markets where players arrive stochastically over time and are drawn from a continuum of types. The cost of matching a client and provider varies, so a social planner is faced with two contending objectives: *a)* to reduce players' *waiting time* before getting matched; and *b)* to form efficient pairs in order to reduce *matching costs*. We show that such markets are characterized by a *quick or cheap* dilemma.

Under a large class of distributional assumptions, there is no 'free lunch', i.e., there exists no clearing schedule that is approximately optimal in terms of both waiting time and matching cost. Building on the no free lunch result, we proceed to fill the spectrum between matching cost and waiting time minimization. We identify a unique breaking point signifying a stark reduction in matching cost contrasted by an increase in waiting time. To explore the finer aspects of this trade-off, we introduce a utility model for the social planner whereby the associated utility of matching cost is of the same order as the agents' utility of waiting time. Under this model, we show that there exists a non-trivial clearing schedule achieving this balance, and we show that this schedule is effectively unique (up to asymptotic order considerations).

Finally, we generalize our key findings by studying different decay rates of matching costs (instead of focusing on one decay rate that results from the micro-founded match costs). We identify two regimes. One, where no free lunch continues to hold. The other, where the benefit from waiting is growing quickly enough, such that a window of opportunity opens and it *is* possible to get a free lunch. As before, in both regimes, greedy scheduling is generally sub-optimal.

Compared to the existing literature [1,3-5] on the trade-off between waiting and mismatch (both in economics and computer science) our model introduces *incomplete information* about the distribution of past and future match costs and considers an infinite type space in a tractable model. As a consequence, the social planner tries to resolve the trade-off between matching optimally and waiting time in light of incomplete information. Incomplete information in our setting implies that the social planner must employ clearing schedules that do not take as input the relative strengths of current and future matches (since the latter is unknown), thus yielding qualitatively new results.

In contrast to prior results for markets with one or two types of match costs where lack of information resulted in optimality of some form of greedy scheduling, we find that greedy clearing is generally not optimal in the presence of many types. Hence, the quick-versus-cheap trade-off is more intricate than previously found. Moreover, our results may actually also have consequences for applications that have been studied before too (e.g. kidney exchange) if other match value metrics (e.g. potential years of life lost or disability-adjusted life years) are used that would produce more than binary match values. By studying fully heterogeneous match costs we have to rely on different mathematical tools compared to previous analyses, which were often able to reduce the induced dynamics to discrete Markov processes.

The key technical innovations of our paper concern the concurrent consideration of a continuum of types, independent arrivals, and incomplete information. In turn, these contributions rely on a range of previously unused tools from probability theory and disordered systems to obtain closed-form solutions. These underlying results are concerned with the expected matching cost for given instances of random, static assignment games. In particular, in static assignment games with the same number of clients and providers and $exp(1)$ distributed edge weights, [2] proved the long-standing conjecture that the expected minimum weight matching converges to $\pi^2/6$ (i.e., as the number of players is growing). This result was later extended by [6] to assignment games with match costs drawn from non-identical exponential distributions. Building on this, we are able to compute the expected matching cost for every `snapshot in time' of the dynamic clearing game. This provides strong foundations for our proofs which are then focused on estimating the fluctuations that result from the random arrival of clients and providers and their randomly drawn match costs. To achieve this, we use several approximation techniques (in particular, the approximation of the arrival process by a continuous-time Wiener process), which allow us to port over several results from martingale limit theory (such as the law of the iterated logarithm).

CCS Concepts: • **Mathematics of computing → Stochastic processes**; *Matching*; Online markets;

KEYWORDS
ACM proceedings; dynamic matching, online markets, market design, assignment game, online matching

REFERENCES
[1] M. Akbarpour, L. Shengwu, and S. Oveis Gharan. 2020. Thickness and information in dynamic matching markets. Journal of Political Economy, 128.3: 783-815.
[2] D. J. Aldous. 2001. The ζ(2) limit in the random assignment problem. Random Structures Algorithms, 18, 381-418.
[3] I. Ashlagi, M. Burq, P. Jaillet, and V. Manshadi, V. 2019. On matching and thickness in heterogeneous dynamic markets. Operations Research, 67(4), 927-949.
[4] I. Ashlagi, A. Nikzad, and P. Strack. 2019. Matching in dynamic imbalanced markets. Working paper. https://arxiv.org/abs/1809.06824
[5] M. Baccara, S. Lee, and L. Yariv. 2020. Optimal dynamic matching. Theoretical Economics, forthcoming.
[6] J. Waestlund. 2005. A proof of a conjecture of Buck, Chan, and Robbins on the expected value of the minimum assignment. Random Structures Algorithms, 26, 237-251.s