Asymptotic welfare performance of Boston assignment algorithms

Mark C. Wilson (https://markcwilson.site)
joint work with Geoffrey Pritchard
accepted by *Stochastic Systems*, 2022

CMSS Workshop, 2022-12-20

Formal model

- A: finite set of n agents; O: set of n objects.

Formal model

- A: finite set of n agents; O: set of n objects.
- Set of all strict linear orders of objects: $L(O)$.

Formal model

- A: finite set of n agents; O: set of n objects.
- Set of all strict linear orders of objects: $L(O)$.
- Set of all profiles is $X := L(O)^A$.

Formal model

- A: finite set of n agents; O: set of n objects.
- Set of all strict linear orders of objects: $L(O)$.
- Set of all profiles is $X := L(O)^A$.
- A matching is a bijection $A \to O$; the set of all such is $M(A, O)$.

Formal model

- A: finite set of n agents; O: set of n objects.
- Set of all strict linear orders of objects: $L(O)$.
- Set of all profiles is $X := L(O)^A$.
- A matching is a bijection $A \to O$; the set of all such is $M(A, O)$.
- The house allocation problem: find a matching rule $f : X \to M(A, O)$.

Informal description

- Given strict ordinal preferences of agents over objects, match each agent with an object!

Informal description

- Given strict ordinal preferences of agents over objects, match each agent with an object!
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.

Informal description

- Given strict ordinal preferences of agents over objects, match each agent with an object!
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Closely related problems: school choice, multi-unit assignment.

Informal description

- Given strict ordinal preferences of agents over objects, match each agent with an object!
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Closely related problems: school choice, multi-unit assignment.
- Key standard axiomatic properties:

Informal description

- Given strict ordinal preferences of agents over objects, match each agent with an object!
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Closely related problems: school choice, multi-unit assignment.
- Key standard axiomatic properties:
  - Pareto efficiency: can't help someone without hurting someone else;

Informal description

- Given strict ordinal preferences of agents over objects, match each agent with an object!
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Closely related problems: school choice, multi-unit assignment.
- Key standard axiomatic properties:
  - Pareto efficiency: can't help someone without hurting someone else;
  - Strategyproofness: no agent ever has incentive to lie about their preferences.

Commonly used solution: serial dictatorship (SD)

- Fix an exogenous order on agents.

Commonly used solution: serial dictatorship (SD)

- Fix an exogenous order on agents.
- Let them choose in turn, according to this order, their favorite remaining object.

Commonly used solution: serial dictatorship (SD)

- Fix an exogenous order on agents.
- Let them choose in turn, according to this order, their favorite remaining object.

Commonly used solution: serial dictatorship (SD)

- Fix an exogenous order on agents.
- Let them choose in turn, according to this order, their favorite remaining object.

This rule is strategyproof, Pareto efficient and easy to implement.

Another solution: Naive Boston

- Choose an exogenous tiebreaking order.

Another solution: Naive Boston

- Choose an exogenous tiebreaking order.
- In round i, all remaining agents bid for their ith preference.

Another solution: Naive Boston

- Choose an exogenous tiebreaking order.
- In round i, all remaining agents bid for their ith preference.
- Use the tiebreaking order to decide who gets an object.

Another solution: Naive Boston

- Choose an exogenous tiebreaking order.
- In round i, all remaining agents bid for their ith preference.
- Use the tiebreaking order to decide who gets an object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

Another solution: Naive Boston

- Choose an exogenous tiebreaking order.
- In round i, all remaining agents bid for their ith preference.
- Use the tiebreaking order to decide who gets an object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

Another solution: Naive Boston

- Choose an exogenous tiebreaking order.
- In round i, all remaining agents bid for their ith preference.
- Use the tiebreaking order to decide who gets an object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

This rule is not strategyproof, but is Pareto efficient and easy to implement.

Adaptive Boston

- Choose an exogenous tiebreaking order.

Adaptive Boston

- Choose an exogenous tiebreaking order.
- In round i, all remaining agents bid for their most preferred remaining object.

Adaptive Boston

- Choose an exogenous tiebreaking order.
- In round i, all remaining agents bid for their most preferred remaining object.
- Use the tiebreaking order to decide who gets an object.

Adaptive Boston

- Choose an exogenous tiebreaking order.
- In round i, all remaining agents bid for their most preferred remaining object.
- Use the tiebreaking order to decide who gets an object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

Adaptive Boston

- Choose an exogenous tiebreaking order.
- In round i, all remaining agents bid for their most preferred remaining object.
- Use the tiebreaking order to decide who gets an object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

Adaptive Boston

- Choose an exogenous tiebreaking order.
- In round i, all remaining agents bid for their most preferred remaining object.
- Use the tiebreaking order to decide who gets an object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

This rule is not strategyproof, but is Pareto efficient and easy to implement.

Preference model

- We consider large random markets with heterogeneous preferences.

Preference model

- We consider large random markets with heterogeneous preferences.
- More precisely, we use the Impartial Culture, where agents sample independently from the uniform distribution on preference orders.

Preference model

- We consider large random markets with heterogeneous preferences.
- More precisely, we use the Impartial Culture, where agents sample independently from the uniform distribution on preference orders.
- This allows for very precise limiting results on performance of our algorithms.

Preference model

- We consider large random markets with heterogeneous preferences.
- More precisely, we use the Impartial Culture, where agents sample independently from the uniform distribution on preference orders.
- This allows for very precise limiting results on performance of our algorithms.
- The Boston algorithms have a balls-in-bins interpretation: each bid is a ball, and each object is a bin.

Assumptions on agent behavior

- We consider only sincere play by agents.

Assumptions on agent behavior

- We consider only sincere play by agents.
- Justification 1: if each agent knows only the (uniform) distribution of preferences of the other agents, then there is no incentive to deviate from sincerity; in other words, the sincere profile is a Bayesian Nash equilibrium for the game induced by the mechanism.

Assumptions on agent behavior

- We consider only sincere play by agents.
- Justification 1: if each agent knows only the (uniform) distribution of preferences of the other agents, then there is no incentive to deviate from sincerity; in other words, the sincere profile is a Bayesian Nash equilibrium for the game induced by the mechanism.
- Justification 2: understanding the social choice rule underlying a given mechanism is a first important step when comparing mechanisms, and the approach is often used in the literature.

Assumptions on agent behavior

- We consider only sincere play by agents.
- Justification 1: if each agent knows only the (uniform) distribution of preferences of the other agents, then there is no incentive to deviate from sincerity; in other words, the sincere profile is a Bayesian Nash equilibrium for the game induced by the mechanism.
- Justification 2: understanding the social choice rule underlying a given mechanism is a first important step when comparing mechanisms, and the approach is often used in the literature.

Assumptions on agent behavior

- We consider only sincere play by agents.
- Justification 1: if each agent knows only the (uniform) distribution of preferences of the other agents, then there is no incentive to deviate from sincerity; in other words, the sincere profile is a Bayesian Nash equilibrium for the game induced by the mechanism.
- Justification 2: understanding the social choice rule underlying a given mechanism is a first important step when comparing mechanisms, and the approach is often used in the literature.

In any case, from now on we shall ignore any issues of strategic behavior by agents.

Overview of key results for Boston

- With high probability most of the agents can be matched to one of their first few preferences — and both Boston algorithms will successfully do so.

Overview of key results for Boston

- With high probability most of the agents can be matched to one of their first few preferences — and both Boston algorithms will successfully do so.
- In particular, the value of the preference rank R an agent is likely to obtain has much the same distribution for any sufficiently large n.

Overview of key results for Boston

- With high probability most of the agents can be matched to one of their first few preferences — and both Boston algorithms will successfully do so.
- In particular, the value of the preference rank R an agent is likely to obtain has much the same distribution for any sufficiently large n.
- More precisely, a given agent's preference rank R for the item he is assigned has a probability distribution that converges to a limit as $n \to \infty$.

Overview of key results for Boston

- With high probability most of the agents can be matched to one of their first few preferences — and both Boston algorithms will successfully do so.
- In particular, the value of the preference rank R an agent is likely to obtain has much the same distribution for any sufficiently large n.
- More precisely, a given agent's preference rank R for the item he is assigned has a probability distribution that converges to a limit as $n \to \infty$.
- We describe explicitly the limiting distributions for the preference rank $R(\theta)$ obtained by an agent in position $\theta \in [0, 1]$ in the choosing order, as a function of θ.

Overview of key results for Boston

- With high probability most of the agents can be matched to one of their first few preferences — and both Boston algorithms will successfully do so.
- In particular, the value of the preference rank R an agent is likely to obtain has much the same distribution for any sufficiently large n.
- More precisely, a given agent's preference rank R for the item he is assigned has a probability distribution that converges to a limit as $n \to \infty$.
- We describe explicitly the limiting distributions for the preference rank $R(\theta)$ obtained by an agent in position $\theta \in [0, 1]$ in the choosing order, as a function of θ.
- These results do not hold for SD, because of the way the last agents are treated.

Key recursive quantities

Definition

The sequence $(\omega_r)_{r=1}^{\infty}$ is defined by $\omega_1 = 1$ and the recursion $\omega_{r+1} = \omega_r e^{-\omega_r}$ for $r \geq 1$.

Lemma

For all $r \geq 3$,

$$\frac{1}{r + \log r} < \omega_r < \frac{1}{r}.$$

Theorem (Number of agents remaining)

Fix $r \geq 1$ and a relative position $\theta \in [0,1]$. The number $N_n(r,\theta)$ of agents starting at position $\leq \theta$ who are still present at round r satisfies

$$\frac{1}{n} N_n(r,\theta) \xrightarrow{p} z_r(\theta)$$

where $z_1(\theta) = \theta$ and

$$z_{r+1}(\theta) = z_r(\theta) - \left(1 - e^{-z_r(\theta)}\right)\omega_r \qquad \text{for } r \geq 1 \tag{1}$$

for Naive Boston. For Adaptive Boston the analogous result is

$$\frac{1}{n} N_n(r,\theta) \xrightarrow{p} y_r(\theta)$$

where $y_1(\theta) = \theta$ and

$$y_{r+1}(\theta) = y_r(\theta) - e^{1-r}\left(1 - \exp\left(-e^{r-1} y_r(\theta)\right)\right) \qquad \text{for } r \geq 1. \tag{2}$$

Number of agents remaining at a given round

Corollary

the total number of agents (and of items) present at round r satisfies

$$\frac{1}{n} N_n(r, 1) \xrightarrow{p} \omega_r$$

for Naive Boston and

$$\frac{1}{n} N_n(r, 1) \xrightarrow{p} e^{1-r}$$

for Adaptive Boston.

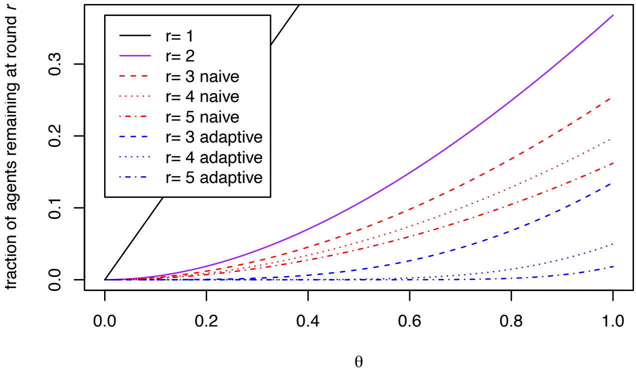

Figure: The limiting fraction of agents with relative position θ or better who survive until round r.

quantity	$r = 1$	$r = 2$	$r = 3$
P(here)	1	$1 - e^{-\theta}$	$(1 - e^{-\theta}) \left(1 - e^{-\theta - e^{-\theta}}\right)$
P(stay | here)	$1 - e^{-\theta}$	$1 - e^{-\theta - e^{-\theta}}$	$1 - e^{-\theta - e^{-\theta} - e^{-\theta - e^{-\theta}}}$

Table: Limiting quantities as $n \to \infty$ for the early rounds of the Naive Boston algorithm for agent at relative position θ.

- This gives us enough information to determine the welfare outcomes, because exiting at round r means the agent gets his rth preference.

quantity	$r = 1$	$r = 2$	$r = 3$
P(here)	1	$1 - e^{-\theta}$	$(1 - e^{-\theta})\left(1 - e^{-\theta - e^{-\theta}}\right)$
P(stay | here)	$1 - e^{-\theta}$	$1 - e^{-\theta - e^{-\theta}}$	$1 - e^{-\theta - e^{-\theta} - e^{-\theta - e^{-\theta}}}$

Table: Limiting quantities as $n \to \infty$ for the early rounds of the Naive Boston algorithm for agent at relative position θ.

- This gives us enough information to determine the welfare outcomes, because exiting at round r means the agent gets his rth preference.
- The same is NOT true for Adaptive Boston.

Extra stuff for Adaptive Boston

- An agent who exits at round $r > 1$ can get any rank object s with $s \geq r$.

Extra stuff for Adaptive Boston

- An agent who exits at round $r > 1$ can get any rank object s with $s \geq r$.
- The quantity u_{rs}, the limiting probability that an agent bids for his sth choice given that he is present at round r, is recursively defined and crucial to the analysis.

Extra stuff for Adaptive Boston

- An agent who exits at round $r > 1$ can get any rank object s with $s \geq r$.
- The quantity u_{rs}, the limiting probability that an agent bids for his sth choice given that he is present at round r, is recursively defined and crucial to the analysis.
- As for Naive Boston we derive limiting distributions for the exit time and the probability of exiting at a given round, but also the preference rank bid for in a given round, and the rank of the item obtained.

Welfare measures

- We convert the rank information into welfare using a scoring rule.

Welfare measures

- We convert the rank information into welfare using a scoring rule.
- We consider here only k-approval welfare (probability of getting one of your top k choices) and normalized Borda.

Welfare measures

- We convert the rank information into welfare using a scoring rule.
- We consider here only k-approval welfare (probability of getting one of your top k choices) and normalized Borda.
- To aggregate over agents, we consider utilitarian welfare, the average welfare over agents, and order bias, the difference between welfare of the first and last agent in the tiebreak order.

Welfare measures

- We convert the rank information into welfare using a scoring rule.
- We consider here only k-approval welfare (probability of getting one of your top k choices) and normalized Borda.
- To aggregate over agents, we consider utilitarian welfare, the average welfare over agents, and order bias, the difference between welfare of the first and last agent in the tiebreak order.
- For all three algorithms, the first agent always gets maximum possible welfare, and the last agent gets worst welfare, so computing order bias reduces to looking at the fate of the last agent.

Theorem

The average utilitarian k-approval welfare over all agents satisfies

$$\frac{1}{n}W_n(1) \xrightarrow{p} \begin{cases} 1 - \omega_{k+1} & \text{for Naive Boston} \\ (1-e^{-1}) \sum_{\{(r,s):r \le s \le k\}} e^{1-r} u_{rs} & \text{for Adaptive Boston} \\ \frac{k}{k+1} & \text{for serial dictatorship.} \end{cases}$$

The sequence ω_k is defined by $\omega_1 = 1$ and the recursion $\omega_{k+1} = \omega_k e^{-\omega_k}$ for $k \ge 1$. The bivariate sequence u_{rs} is defined by the recursion $u_{11} = 1$, $u_{1,s} = 0$ for $s > 1$, $u_{rs} = 0$ for $s < r$, and

$$u_{rs} = e^{1-r} u_{r-1,s-1} + (1-e^{1-r}) u_{r,s-1}$$

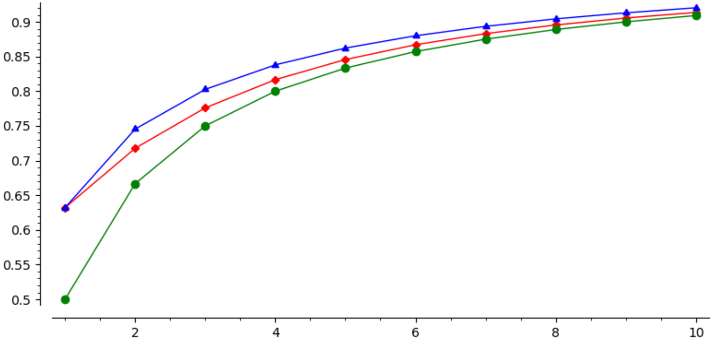

Figure: Limiting values as $n \to \infty$ of k-approval welfare, for $1 \leq k \leq 10$. Bottom: Serial Dictatorship. Middle: Adaptive Boston. Top: Naive Boston.

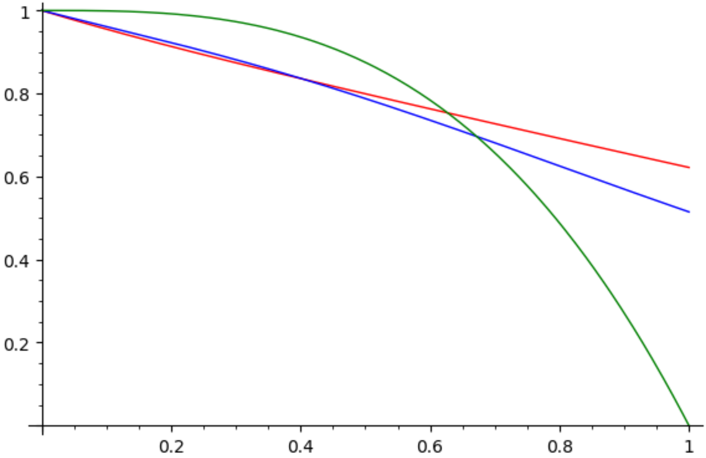

Figure: Limiting values as $n \to \infty$ of the probability of getting at worst one's 3rd choice, as a function of θ. Serial Dictatorship. Adaptive Boston. Naive Boston.

Theorem

For each fixed k, the k-approval order bias of Naive Boston is asymptotically $\prod_{j=2}^{k+1}(1-\omega_j)$, while for Adaptive Boston it is asymptotically

$$1 - e^{-1} \sum_{\{(r,s):r\leq s\leq k\}} \left(1-e^{-1}\right)^{r-1} u_{rs}.$$

and for Serial Dictatorship it is asymptotically 1.

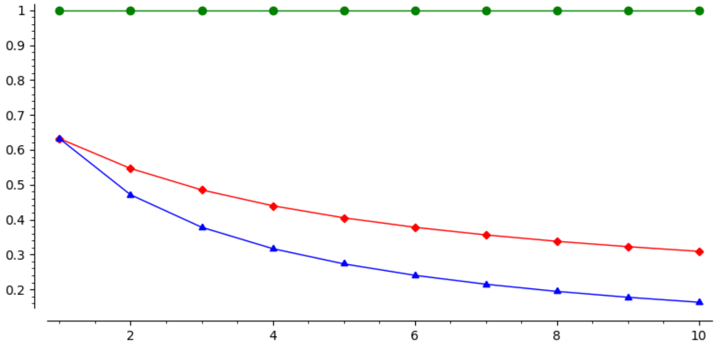

Figure: Limiting values as $n \to \infty$ of k-approval order bias, for $1 \leq k \leq 10$. Top: Serial Dictatorship. Middle: Adaptive Boston. Bottom: Naive Boston.

Differences between SD and Boston

- There is no limiting distribution for SD. For example, the preference rank of the last agent's object is uniformly distributed over all n objects.

Differences between SD and Boston

- There is no limiting distribution for SD. For example, the preference rank of the last agent's object is uniformly distributed over all n objects.
- The last 20% of agents under SD do much worse in expectation than under the Boston algorithms.

Differences between SD and Boston

- There is no limiting distribution for SD. For example, the preference rank of the last agent's object is uniformly distributed over all n objects.
- The last 20% of agents under SD do much worse in expectation than under the Boston algorithms.
- Naive Boston provably beats SD on welfare and order bias, and we suspect that Adaptive Boston does too.

Unsolved questions

- What is the expected rank of the item gained by a random agent?

Unsolved questions

- What is the expected rank of the item gained by a random agent?
- For SD, the asymptotic answer $\Theta(\log n)$ was derived by Frieze & Pittel and this was refined to an exact formula $((n+1)H_n - n)/n \sim \log n$ by Knuth.

Unsolved questions

- What is the expected rank of the item gained by a random agent?
- For SD, the asymptotic answer $\Theta(\log n)$ was derived by Frieze & Pittel and this was refined to an exact formula $((n+1)H_n - n)/n \sim \log n$ by Knuth.
- For the Boston mechanisms, our results show only that the rank is $o(n)$, but we suspect that it is $\Theta(\log n)$.

Unsolved questions

- What is the expected rank of the item gained by a random agent?
- For SD, the asymptotic answer $\Theta(\log n)$ was derived by Frieze & Pittel and this was refined to an exact formula $((n+1)H_n - n)/n \sim \log n$ by Knuth.
- For the Boston mechanisms, our results show only that the rank is $o(n)$, but we suspect that it is $\Theta(\log n)$.
- It is known that for the rank-maximizing mechanism RM, which maximizes the number of agents receiving their first choice, then subject to that the number of agents receiving their second choice, etc, the expected average rank in a random market is asymptotically constant.

Unsolved questions

- What is the expected rank of the item gained by a random agent?
- For SD, the asymptotic answer $\Theta(\log n)$ was derived by Frieze & Pittel and this was refined to an exact formula $((n+1)H_n - n)/n \sim \log n$ by Knuth.
- For the Boston mechanisms, our results show only that the rank is $o(n)$, but we suspect that it is $\Theta(\log n)$.
- It is known that for the rank-maximizing mechanism RM, which maximizes the number of agents receiving their first choice, then subject to that the number of agents receiving their second choice, etc, the expected average rank in a random market is asymptotically constant.
- Although similar to RM at first sight, Naive Boston is not as strict, since it makes a choice based on tiebreaking at the first round, and hence may diverge from RM even at the second round.

Possible extensions

- The Boston algorithms discussed here are specializations of algorithms commonly used for school choice to the case where each school has a single seat and schools have a common preference order over applicants.

Possible extensions

- The Boston algorithms discussed here are specializations of algorithms commonly used for school choice to the case where each school has a single seat and schools have a common preference order over applicants.
- Further analysis in the general school choice case would be very desirable.

Possible extensions

- The Boston algorithms discussed here are specializations of algorithms commonly used for school choice to the case where each school has a single seat and schools have a common preference order over applicants.
- Further analysis in the general school choice case would be very desirable.
- It would also be interesting to study welfare and order bias in the multi-unit assignment model.

Possible extensions

- The Boston algorithms discussed here are specializations of algorithms commonly used for school choice to the case where each school has a single seat and schools have a common preference order over applicants.
- Further analysis in the general school choice case would be very desirable.
- It would also be interesting to study welfare and order bias in the multi-unit assignment model.
- Simulation shows that the Mallows preference model yields the same ranking of welfare performance of the algorithms.