Dynamic Matching and Learning

Yilmaz Kocer

USC

NEGT 2014, Caltech

Motivation

- Dynamic matching

Motivation

- Dynamic matching
- Incomplete Information

Motivation

- Dynamic matching
- Incomplete Information
- No transfers

Motivation

- Dynamic matching
- Incomplete Information
- No transfers
- Strategy of the matched agents or stability is not the focus; the focus is on the information mediating mechanism.

Motivation

- Dynamic matching
- Incomplete Information
- No transfers
- Strategy of the matched agents or stability is not the focus; the focus is on the information mediating mechanism.
- Combinatorial aspects of matching and learning

Motivation

- Dynamic matching
- Incomplete Information
- No transfers
- Strategy of the matched agents or stability is not the focus; the focus is on the information mediating mechanism.
- Combinatorial aspects of matching and learning
- example: Online dating, user content or ads on screen vs user types...

Desiderata

- Build a theory of matching with;

Desiderata

- Build a theory of matching with;
- Incomplete Information about own preferences, partner's type.

Desiderata

- Build a theory of matching with;
- Incomplete Information about own preferences, partner's type.
- Dynamics, learning: information is to be had over time, through matches.

Desiderata

- Build a theory of matching with;
- Incomplete Information about own preferences, partner's type.
- Dynamics, learning: information is to be had over time, through matches.
- Central mechanism as an facilitator of increased match quality.

Literature

- Matching with Incomplete Information, Liu et al 2014, Bikchandani 2014

Literature

- Matching with Incomplete Information, Liu et al 2014, Bikchandani 2014
- Dynamic Mechanism Design, Bergemann &Valimaki 2010, Athey & Segal 2013, Kakade Lobel & Nazerzadeh 2011

Literature

- Matching with Incomplete Information, Liu et al 2014, Bikchandani 2014
- Dynamic Mechanism Design, Bergemann &Valimaki 2010, Athey & Segal 2013, Kakade Lobel & Nazerzadeh 2011
- Labor search with a "matching function", Anderson & Smith 2010 , Shimer & Smith 2000 assortative matching...

Literature

- Matching with Incomplete Information, Liu et al 2014, Bikchandani 2014
- Dynamic Mechanism Design, Bergemann &Valimaki 2010, Athey & Segal 2013, Kakade Lobel & Nazerzadeh 2011
- Labor search with a "matching function", Anderson & Smith 2010 , Shimer & Smith 2000 assortative matching...
- Matching with costly signaling or interviews, Lee & Schwartz 2012, Hoppe 2009 Nahum et al 2013

Literature

- Matching with Incomplete Information, Liu et al 2014, Bikchandani 2014
- Dynamic Mechanism Design, Bergemann &Valimaki 2010, Athey & Segal 2013, Kakade Lobel & Nazerzadeh 2011
- Labor search with a "matching function", Anderson & Smith 2010 , Shimer & Smith 2000 assortative matching...
- Matching with costly signaling or interviews, Lee & Schwartz 2012, Hoppe 2009 Nahum et al 2013
- One sided matching Bloch & Cantala 2013, 2014, Gershkov Moldovanu 2010, 2012

Literature

- Matching with Incomplete Information, Liu et al 2014, Bikchandani 2014
- Dynamic Mechanism Design, Bergemann &Valimaki 2010, Athey & Segal 2013, Kakade Lobel & Nazerzadeh 2011
- Labor search with a "matching function", Anderson & Smith 2010 , Shimer & Smith 2000 assortative matching...
- Matching with costly signaling or interviews, Lee & Schwartz 2012, Hoppe 2009 Nahum et al 2013
- One sided matching Bloch & Cantala 2013, 2014, Gershkov Moldovanu 2010, 2012
- Combinatorial bandits - regret bounds, Chen, Wang & Yuan 2013

Literature

- Matching with Incomplete Information, Liu et al 2014, Bikchandani 2014
- Dynamic Mechanism Design, Bergemann &Valimaki 2010, Athey & Segal 2013, Kakade Lobel & Nazerzadeh 2011
- Labor search with a "matching function", Anderson & Smith 2010 , Shimer & Smith 2000 assortative matching...
- Matching with costly signaling or interviews, Lee & Schwartz 2012, Hoppe 2009 Nahum et al 2013
- One sided matching Bloch & Cantala 2013, 2014, Gershkov Moldovanu 2010, 2012
- Combinatorial bandits - regret bounds, Chen, Wang & Yuan 2013
- Two sided matching with learning: Rastegari et al 2013 , Das & Kamenica 2008,

Literature

- Matching with Incomplete Information, Liu et al 2014, Bikchandani 2014
- Dynamic Mechanism Design, Bergemann &Valimaki 2010, Athey & Segal 2013, Kakade Lobel & Nazerzadeh 2011
- Labor search with a "matching function", Anderson & Smith 2010 , Shimer & Smith 2000 assortative matching...
- Matching with costly signaling or interviews, Lee & Schwartz 2012, Hoppe 2009 Nahum et al 2013
- One sided matching Bloch & Cantala 2013, 2014, Gershkov Moldovanu 2010, 2012
- Combinatorial bandits - regret bounds, Chen, Wang & Yuan 2013
- Two sided matching with learning: Rastegari et al 2013 , Das & Kamenica 2008,
- Matching with complete info, Kennes Monte 2014 Pereyra 2013

Literature

- Matching with Incomplete Information, Liu et al 2014, Bikchandani 2014
- Dynamic Mechanism Design, Bergemann &Valimaki 2010, Athey & Segal 2013, Kakade Lobel & Nazerzadeh 2011
- Labor search with a "matching function", Anderson & Smith 2010 , Shimer & Smith 2000 assortative matching...
- Matching with costly signaling or interviews, Lee & Schwartz 2012, Hoppe 2009 Nahum et al 2013
- One sided matching Bloch & Cantala 2013, 2014, Gershkov Moldovanu 2010, 2012
- Combinatorial bandits - regret bounds, Chen, Wang & Yuan 2013
- Two sided matching with learning: Rastegari et al 2013 , Das & Kamenica 2008,
- Matching with complete info, Kennes Monte 2014 Pereyra 2013
- Strategic matching aligned preferences Niederle &Yariv 2009 ....

Dynamic Matching as Combinatorial Bandits

Consider dynamically matching 2 men and 2 women to maximize the expected $\beta-$ discounted sum of match payoffs up to infinity, equivalently the match payoff at a geometric stopping time with $1-\beta$. Each **(man i,woman j)** is an *independent* bandit, whereas the overall match permutation is not:

Payoffs to a matching pair (man i, woman j) each period	Woman 1	Woman 2
man 1	1+e	1/2 prob. → 2 1/2 prob. → 0
man 2	1/2 prob. → 2 1/2 prob. → 0	1+e

Dynamic Matching as Combinatorial Bandits

- Consider each bandit in isolation. Remember the Gittins index for a general bandit arm with state evolution $\{x_t\}_{t=1}^{\infty}$ at state x_0 is

Dynamic Matching as Combinatorial Bandits

- Consider each bandit in isolation. Remember the Gittins index for a general bandit arm with state evolution $\{x_t\}_{t=1}^{\infty}$ at state x_0 is
-
$$\mathbb{G}(x_0) = \sup_{\tau>0} \frac{\mathbb{E}\{\Sigma_{t=0}^{t=\tau-1} \beta^t u(x_t)\}}{\mathbb{E}\{\Sigma_{t=0}^{t=\tau-1} \beta^t\}}$$

  for all stopping times τ

Dynamic Matching as Combinatorial Bandits

- Consider each bandit in isolation. Remember the Gittins index for a general bandit arm with state evolution $\{x_t\}_{t=1}^{\infty}$ at state x_0 is
- $$\mathbb{G}(x_0) = \sup_{\tau>0} \frac{\mathbb{E}\{\Sigma_{t=0}^{t=\tau-1} \beta^t u(x_t)\}}{\mathbb{E}\{\Sigma_{t=0}^{t=\tau-1} \beta^t\}}$$

  for all stopping times τ
- For our example, $\mathbb{G}_{11} = \mathbb{G}_{22} = \frac{1+e}{1-\beta}$ whereas $\mathbb{G}_{12} = \mathbb{G}_{21} = g$

Dynamic Matching as Combinatorial Bandits

- Consider each bandit in isolation. Remember the Gittins index for a general bandit arm with state evolution $\{x_t\}_{t=1}^{\infty}$ at state x_0 is
- $$\mathbb{G}(x_0) = \sup_{\tau>0} \frac{\mathbb{E}\{\Sigma_{t=0}^{t=\tau-1} \beta^t u(x_t)\}}{\mathbb{E}\{\Sigma_{t=0}^{t=\tau-1} \beta^t\}}$$

  for all stopping times τ
- For our example, $\mathbb{G}_{11} = \mathbb{G}_{22} = \frac{1+e}{1-\beta}$ whereas $\mathbb{G}_{12} = \mathbb{G}_{21} = g$
- $$g = \frac{1}{2}\frac{2}{1-\beta} + \frac{1}{2}\beta \Rightarrow g = \frac{2}{(2-\beta)(1-\beta)}$$

Dynamic Matching as Combinatorial Bandits

- As the problem is an MDP, the only two matching strategies to compare are
  (I) "always match $(1,2) \rightarrow (1,2)$"
  (II)"match $(1,2) \rightarrow (2,1)$ first, and then revert to $(1,2) \rightarrow (1,2)$ if not successful"

Dynamic Matching as Combinatorial Bandits

- As the problem is an MDP, the only two matching strategies to compare are
  (I) "always match $(1,2) \rightarrow (1,2)$"
  (II)"match $(1,2) \rightarrow (2,1)$ first, and then revert to $(1,2) \rightarrow (1,2)$ if not successful"
- For $\mathbb{G}_{12} = g > \frac{1+e}{1-\beta}$ slightly, each agent would prefer (II) individually. Socially, one might think going with the strategy that maximizes the sum of the Gittins indices would be better. However,

Dynamic Matching as Combinatorial Bandits

- As the problem is an MDP, the only two matching strategies to compare are
  (I) "always match $(1,2) \to (1,2)$"
  (II)"match $(1,2) \to (2,1)$ first, and then revert to $(1,2) \to (1,2)$ if not successful"
- For $\mathbb{G}_{12} = g > \frac{1+e}{1-\beta}$ slightly, each agent would prefer (II) individually. Socially, one might think going with the strategy that maximizes the sum of the Gittins indices would be better. However,
- The total social welfare for (I) $= 2\mathbb{G}_{11}$

Dynamic Matching as Combinatorial Bandits

- As the problem is an MDP, the only two matching strategies to compare are
  (I) "always match $(1,2) \to (1,2)$"
  (II)"match $(1,2) \to (2,1)$ first, and then revert to $(1,2) \to (1,2)$ if not successful"
- For $\mathbb{G}_{12} = g > \frac{1+e}{1-\beta}$ slightly, each agent would prefer (II) individually. Socially, one might think going with the strategy that maximizes the sum of the Gittins indices would be better. However,
- The total social welfare for (I) $= 2\mathbb{G}_{11}$
-

$$(II) = \frac{1}{4}\frac{4}{1-\beta} + \frac{1}{4}\frac{2(1+e)\beta}{1-\beta} + \frac{1}{2}\{2 + \frac{\beta}{1-\beta} max\{2, 2(1+e)\}$$

For $e \in (\frac{\beta}{4-\beta}, \frac{\beta}{2-\beta})$, $(I) > (II)$!!

Dynamic Matching as Combinatorial Bandits

- Intuitively, in (II) the "sum" of option values of matching $1 \rightarrow 2$ AND $2 \rightarrow 1$ do not materialize. The "Gittins index" of the matching is LESS THAN $2\mathbb{G}_{12}$, as we cannot preferentially stop one leg of the matching keeping the other leg intact.

Dynamic Matching as Combinatorial Bandits

- Intuitively, in (II) the "sum" of option values of matching $1 \to 2$ AND $2 \to 1$ do not materialize. The "Gittins index" of the matching is LESS THAN $2\mathbb{G}_{12}$, as we cannot preferentially stop one leg of the matching keeping the other leg intact.
- It cannot be remedied even if the agents come from a pool of 2 types, and we can match two men to two women of the same type, as long as we match 2 pairs simultaneously.

Dynamic Matching as Combinatorial Bandits

- Intuitively, in (II) the "sum" of option values of matching $1 \to 2$ AND $2 \to 1$ do not materialize. The "Gittins index" of the matching is LESS THAN $2\mathbb{G}_{12}$, as we cannot preferentially stop one leg of the matching keeping the other leg intact.
- It cannot be remedied even if the agents come from a pool of 2 types, and we can match two men to two women of the same type, as long as we match 2 pairs simultaneously.
- One way to go around it to give up the simultaneity of the whole matching, and match agents one by one.

Dynamic Matching as Combinatorial Bandits

- Intuitively, in (II) the "sum" of option values of matching $1 \to 2$ AND $2 \to 1$ do not materialize. The "Gittins index" of the matching is LESS THAN $2\mathbb{G}_{12}$, as we cannot preferentially stop one leg of the matching keeping the other leg intact.
- It cannot be remedied even if the agents come from a pool of 2 types, and we can match two men to two women of the same type, as long as we match 2 pairs simultaneously.
- One way to go around it to give up the simultaneity of the whole matching, and match agents one by one.
- Another way is to restrict the space of preferences allowed ; luckily, for some strongly correlated preferences, we will retain the indexability.

- There are N men and N women in the match pool, they are infinitely lived.

MODEL Agents

- There are N men and N women in the match pool, they are infinitely lived.
- Preferences of the agents are not known ex-ante: There is an 'optimal matching", a permutation $\bar{\sigma} : N \to N$, that is unknown to all parties. That is, it is common knowledge that each type has a single favorite type on the opposite market, and that each type's favorite is distinct.

- There are N men and N women in the match pool, they are infinitely lived.
- Preferences of the agents are not known ex-ante: There is an 'optimal matching", a permutation $\bar{\sigma} : N \to N$, that is unknown to all parties. That is, it is common knowledge that each type has a single favorite type on the opposite market, and that each type's favorite is distinct.
- Nature chooses $\bar{\sigma}$ according to the common prior $\pi : N \times N \to \mathbb{R}_+$ with row and column sums 1; a doubly-stochastic matrix.

MODEL Agents

- There are N men and N women in the match pool, they are infinitely lived.
- Preferences of the agents are not known ex-ante: There is an 'optimal matching", a permutation $\bar{\sigma} : N \to N$, that is unknown to all parties. That is, it is common knowledge that each type has a single favorite type on the opposite market, and that each type's favorite is distinct.
- Nature chooses $\bar{\sigma}$ according to the common prior $\pi : N \times N \to \mathbb{R}_+$ with row and column sums 1; a doubly-stochastic matrix.
- Notice that $\pi(i, j)$ is the *marginal* probability that man i finds a woman j a good match, **and** vice versa.

Timing Reports and Rematch

- Time is discrete $t = 0, 1, 2, \ldots$

Timing Reports and Rematch

- Time is discrete $t = 0, 1, 2, ...$
- There is a central mechanism (CM), matching each man with exactly one distinct woman each period; $\sigma_t : N \to N$, a permutation on N, with Σ the set of all permutations on types (all matchings).

Timing Reports and Rematch

- Time is discrete $t = 0, 1, 2, \ldots$
- There is a central mechanism (CM), matching each man with exactly one distinct woman each period; $\sigma_t : N \to N$, a permutation on N, with Σ the set of all permutations on types (all matchings).
- Each period, each matched pair (i, j) receives a match quality signal $s \in \{1 = \mathit{success}, 0 = \mathit{fail}\}$ with probability $a > 1/2$ if $\bar{\sigma}(i) = j$ and with probability $1 - a$ otherwise.

- The signals are observed by CM, but not by other pairs. Notice that one match's signal is informative for other matches, too.

- The signals are observed by CM, but not by other pairs. Notice that one match's signal is informative for other matches, too.
- After receiving the signal, each agent privately declares $d = 1 = \mathit{Accept}$ or $d = 0 = \mathit{Reject}$ for the current partner; acceptance is irrevocable.

- The signals are observed by CM, but not by other pairs. Notice that one match's signal is informative for other matches, too.
- After receiving the signal, each agent privately declares $d = 1 = Accept$ or $d = 0 = Reject$ for the current partner; acceptance is irrevocable.
- With an exogenous probability $1 - \beta$, the dating stops at the end of each dating period.

- The signals are observed by CM, but not by other pairs. Notice that one match's signal is informative for other matches, too.
- After receiving the signal, each agent privately declares $d = 1 = Accept$ or $d = 0 = Reject$ for the current partner; acceptance is irrevocable.
- With an exogenous probability $1 - \beta$, the dating stops at the end of each dating period.
- If there are reciprocally accepted matches they end up as a final match, if not, they are left single.

- Agents have an outside option of 0, and receive payoffs of $\beta^t \times 1$ iff the dating stops at time t with an reciprocally accepted *optimal partner* j , i.e. $\bar{\sigma}_t(i) = j$ and 0 otherwise.

- Agents have an outside option of 0, and receive payoffs of $\beta^t \times 1$ iff the dating stops at time t with an reciprocally accepted *optimal partner* j , i.e. $\bar{\sigma}_t(i) = j$ and 0 otherwise.
- Man i 's observed history is, $h_t^i = h_{t-1}^i \oplus (\sigma_t(i), s^i_{\sigma_t(i),t}, d_t)$ with $h_0^i = \emptyset$ and the CM's observed history is $h_t = \oplus_i h_t^i \oplus_j h_t^j$; and capitals H_t^i, H_t for the *set* of histories.

- Agents have an outside option of 0, and receive payoffs of $\beta^t \times 1$ iff the dating stops at time t with an reciprocally accepted *optimal partner* j , i.e. $\bar{\sigma}_t(i) = j$ and 0 otherwise.
- Man i 's observed history is, $h^i_t = h^i_{t-1} \oplus (\sigma_t(i), s^i_{\sigma_t(i),t}, d_t)$ with $h^i_0 = \emptyset$ and the CM's observed history is $h_t = \oplus_i h^i_t \oplus_j h^j_t$; and capitals H^i_t, H_t for the *set* of histories.
- A strategy of an agent i, d^i is a collection of $d^i_t : H^i_t \to \{0, 1\}$, and a matching mechanism $\mathbb{M}$ is a collection of $\mathbb{M}_t : H_t \to \Sigma$

- Agents have an outside option of 0, and receive payoffs of $\beta^t \times 1$ iff the dating stops at time t with an reciprocally accepted *optimal partner* j , i.e. $\bar{\sigma}_t(i) = j$ and 0 otherwise.
- Man i 's observed history is, $h_t^i = h_{t-1}^i \oplus (\sigma_t(i), s^i_{\sigma_t(i),t}, d_t)$ with $h_0^i = \emptyset$ and the CM's observed history is $h_t = \oplus_i h_t^i \oplus_j h_t^j$; and capitals H_t^i, H_t for the *set* of histories.
- A strategy of an agent i, d^i is a collection of $d_t^i : H_t^i \to \{0, 1\}$, and a matching mechanism $\mathbb{M}$ is a collection of $\mathbb{M}_t : H_t \to \Sigma$
- Note that only the first acceptance counts; hence we can focus on histories where that agent hasn't accepted anybody yet; the timing of acceptance is the strategic component.

- I restrict attention to stationary Markov strategies in private beliefs, and stationary Markov Mechanisms: The agent i's posterior belief about the "optimal matching" $\bar{\sigma}$ at time t is derived only from the private history h_t^i. Let the posterior $\pi_t^i = \pi \oplus s_t^i$ be the Bayes update after the collection of own private signals.

- I restrict attention to stationary Markov strategies in private beliefs, and stationary Markov Mechanisms: The agent i's posterior belief about the "optimal matching" $\bar{\sigma}$ at time t is derived only from the private history h_t^i. Let the posterior $\pi_t^i = \pi \oplus s_t^i$ be the Bayes update after the collection of own private signals.
- I assume d^i is a function of π_t^i and the current match but NOT on the time index. Notice that if nobody else (including the mechanism) is using t, there exists a stationary Markov best response not depending on t, as the environment is then stationary.

- I restrict attention to stationary Markov strategies in private beliefs, and stationary Markov Mechanisms: The agent i's posterior belief about the "optimal matching" $\bar{\sigma}$ at time t is derived only from the private history h_t^i. Let the posterior $\pi_t^i = \pi \oplus s_t^i$ be the Bayes update after the collection of own private signals.
- I assume d^i is a function of π_t^i and the current match but NOT on the time index. Notice that if nobody else (including the mechanism) is using t, there exists a stationary Markov best response not depending on t, as the environment is then stationary.
- I assume $\mathbb{M}$ depends ONLY on the posterior π_t^M, which is the Bayes update after observing ALL agents' signals.

Lemma

Given any stationary Markov mechanism $\mathbb{M}$, *there exists a stationary Markov equilibrium for the agents. Furthermore, the equilibrium strategy* d^i *induces a Markov stopping time that solves the Gittins equation;*

$$\sup_{\tau>0} \frac{\mathbb{E}\{\Sigma_{t=0}^{t=\tau-1} u(x_t)\}}{\mathbb{E}\{\Sigma_{t=0}^{t=\tau-1} \beta^t\}}$$

Here, τ *is derived through the first acceptance decision in* d^i. *Expectation is wrt to the Markov evolution of private beliefs implied by the mechanism and others' strategy.*

Efficient Mechanism

- Consider the first best scenario. We will describe an efficient/ first best mechanism.

Efficient Mechanism

- Consider the first best scenario. We will describe an efficient/ first best mechanism.
- Let $\mathbb{M}_\pi = \sigma \in \Sigma$ be the matching induced under the posterior π on the agent preferences (on the optimal -unknown- matching $\bar{\sigma}$)

Efficient Mechanism

- Consider the first best scenario. We will describe an efficient/ first best mechanism.
- Let $\mathbb{M}_\pi = \sigma \in \Sigma$ be the matching induced under the posterior π on the agent preferences (on the optimal -unknown- matching $\bar{\sigma}$)

Efficient Mechanism

- Consider the first best scenario. We will describe an efficient/ first best mechanism.
- Let $\mathbb{M}_\pi = \sigma \in \Sigma$ be the matching induced under the posterior π on the agent preferences (on the optimal -unknown- matching $\bar{\sigma}$)

Theorem

Suppose CM is making the acceptance decisions on behalf of the agents. Let $\mathbb{M}^{eff}$ be the mechanism such that

(i) $\mathbb{M}^{eff}_\pi$ maximizes $\{\sum_{i=1}^{N} \pi_{i,\sigma(i)}\}$ over all permutations.

(ii) The acceptance decisions taken on behalf of the agents, d^i, are derived as in the Lemma, given the matching component $\mathbb{M}^{eff}_\pi$

$\mathbb{M}^{eff}$ maximizes the expected sum of payoffs to all agents.

Efficient Mechanism - **Proof**

- Note that the mechanism uses the matching/permutation that is most likely to be the -unknown- optimal matching, $\bar{\sigma}$

- Note that the mechanism uses the matching/permutation that is most likely to be the -unknown- optimal matching, $\bar{\sigma}$
- This means that myopic optimal matching is efficient in our setup.

Efficient Mechanism - **Proof**

- Note that the mechanism uses the matching/permutation that is most likely to be the -unknown- optimal matching, $\bar{\sigma}$
- This means that myopic optimal matching is efficient in our setup.
- It is the largest probability matching in the support of π. Computing it is very easy; it is the static assignment problem.

Efficient Mechanism - **Proof**

- Note that the mechanism uses the matching/permutation that is most likely to be the -unknown- optimal matching, $\bar{\sigma}$
- This means that myopic optimal matching is efficient in our setup.
- It is the largest probability matching in the support of π. Computing it is very easy; it is the static assignment problem.

Efficient Mechanism - **Proof**

- The proof proceeds through these steps;

Efficient Mechanism - **Proof**

- The proof proceeds through these steps;
- CM with acceptance rights faces a MDP, the actions being
  (1) choosing the permutation each period, and
  (2) deciding on behalf of the agents.

Efficient Mechanism - **Proof**

- The proof proceeds through these steps;
- CM with acceptance rights faces a MDP, the actions being
  (1) choosing the permutation each period, and
  (2) deciding on behalf of the agents.
- The second stage should be a Markov stopping time given (1) is Markov. Hence the whole problem is a MDP.

Efficient Mechanism - **Proof**

- The proof proceeds through these steps;
- CM with acceptance rights faces a MDP, the actions being
  (1) choosing the permutation each period, and
  (2) deciding on behalf of the agents.
- The second stage should be a Markov stopping time given (1) is Markov. Hence the whole problem is a MDP.
- Unimprovability principle implies we just need to show the value function is excessive.

Efficient Mechanism - **Proof**

- The proof proceeds through these steps;
- CM with acceptance rights faces a MDP, the actions being
  (1) choosing the permutation each period, and
  (2) deciding on behalf of the agents.
- The second stage should be a Markov stopping time given (1) is Markov. Hence the whole problem is a MDP.
- Unimprovability principle implies we just need to show the value function is excessive.
- An interchange argument: Swapping two components of any given permutaion towards the myopic optimal permutation improves the payoff.

Efficient Mechanism - **Proof**

- The proof proceeds through these steps;
- CM with acceptance rights faces a MDP, the actions being
  (1) choosing the permutation each period, and
  (2) deciding on behalf of the agents.
- The second stage should be a Markov stopping time given (1) is Markov. Hence the whole problem is a MDP.
- Unimprovability principle implies we just need to show the value function is excessive.
- An interchange argument: Swapping two components of any given permutaion towards the myopic optimal permutation improves the payoff.
- The myopic optimality is reminiscent of early results on bandits where arms are prefectly negatively correlated, i.e. Feldman (1962), Rodman(1978).

Efficient Mechanism $\mathbb{M}^{eff}$ and Incentive Compatibility

- For our strategially limited agents, the only IC is when and whom to accept.

Efficient Mechanism $\mathbb{M}^{eff}$ and Incentive Compatibility

- For our strategially limited agents, the only IC is when and whom to accept.
- We will lastly show that the (2) stage of the efficient mechanism $\mathbb{M}^{eff}$, the acceptance decisions of agents (whom and when) can be easily decentralized.

Efficient Mechanism $\mathbb{M}^{eff}$ and Incentive Compatibility

- For our strategially limited agents, the only IC is when and whom to accept.
- We will lastly show that the (2) stage of the efficient mechanism $\mathbb{M}^{eff}$, the acceptance decisions of agents (whom and when) can be easily decentralized.
- Remember that in Theorem (1), the CM took acceptance decisions on behalf of the agents to maximize the SUM of all agents' payoffs.

Efficient Mechanism $\mathbb{M}^{eff}$ and Incentive Compatibility

- For our strategially limited agents, the only IC is when and whom to accept.
- We will lastly show that the (2) stage of the efficient mechanism $\mathbb{M}^{eff}$, the acceptance decisions of agents (whom and when) can be easily decentralized.
- Remember that in Theorem (1), the CM took acceptance decisions on behalf of the agents to maximize the SUM of all agents' payoffs.

Theorem

If the CM uses the mechanism $\mathbb{M}^{eff}$ it is an equilibrium for the agents to use acceptance decisions to imitate the acceptance decisions of the efficient mechanism.

Efficient Mechanism $\mathbb{M}^{eff}$ and Incentive Compatibility - **Proof**

- The proof is by contradiction. Suppose all other agents follow the efficient mechanism's acceptance decisions, and an agent deviates.

Efficient Mechanism $\mathbb{M}^{eff}$ and Incentive Compatibility - **Proof**

- The proof is by contradiction. Suppose all other agents follow the efficient mechanism's acceptance decisions, and an agent deviates.
- The sign of the first order effect of the deviation is the same for the agent's payoffs and the total payoff (of all agents), implied by the positive correlation in preferences.

Efficient Mechanism $\mathbb{M}^{eff}$ and Incentive Compatibility - **Proof**

- The proof is by contradiction. Suppose all other agents follow the efficient mechanism's acceptance decisions, and an agent deviates.
- The sign of the first order effect of the deviation is the same for the agent's payoffs and the total payoff (of all agents), implied by the positive correlation in preferences.
- Corollary $\Rightarrow$ The expected payoff of each agent in the efficient mechanism is increasing in the informativeness of the signal a, and the discount factor β.

Conclusion and Extensions

- I develop a model of dynamic matching where the CM has a facilitator role for information mediation.

Conclusion and Extensions

- I develop a model of dynamic matching where the CM has a facilitator role for information mediation.
- Strategic considerations on part of the agents were not the focus.

Conclusion and Extensions

- I develop a model of dynamic matching where the CM has a facilitator role for information mediation.
- Strategic considerations on part of the agents were not the focus.
- Combinatorial aspects of learning through a matching is emphasized.

Conclusion and Extensions

- I develop a model of dynamic matching where the CM has a facilitator role for information mediation.
- Strategic considerations on part of the agents were not the focus.
- Combinatorial aspects of learning through a matching is emphasized.
- A clean analysis for a rather restricted class of preferences. An original extension of the Gittins indexability to a combinatorial learning problem.

Conclusion and Extensions

- I develop a model of dynamic matching where the CM has a facilitator role for information mediation.
- Strategic considerations on part of the agents were not the focus.
- Combinatorial aspects of learning through a matching is emphasized.
- A clean analysis for a rather restricted class of preferences. An original extension of the Gittins indexability to a combinatorial learning problem.
- Further steps:

Conclusion and Extensions

- I develop a model of dynamic matching where the CM has a facilitator role for information mediation.
- Strategic considerations on part of the agents were not the focus.
- Combinatorial aspects of learning through a matching is emphasized.
- A clean analysis for a rather restricted class of preferences. An original extension of the Gittins indexability to a combinatorial learning problem.
- Further steps:
- **Entry and exit from the match pool.**

Conclusion and Extensions

- I develop a model of dynamic matching where the CM has a facilitator role for information mediation.
- Strategic considerations on part of the agents were not the focus.
- Combinatorial aspects of learning through a matching is emphasized.
- A clean analysis for a rather restricted class of preferences. An original extension of the Gittins indexability to a combinatorial learning problem.
- Further steps:
- **Entry and exit from the match pool.**
- **Richer preferences.**

Conclusion and Extensions

- I develop a model of dynamic matching where the CM has a facilitator role for information mediation.
- Strategic considerations on part of the agents were not the focus.
- Combinatorial aspects of learning through a matching is emphasized.
- A clean analysis for a rather restricted class of preferences. An original extension of the Gittins indexability to a combinatorial learning problem.
- Further steps:
- **Entry and exit from the match pool.**
- **Richer preferences.**
- **More strategic components on part of the agents.**
  Thank you !