

Knocking out the supply and sorting in centralized job markets*

Guillaume Haeringer[†]

Universitat Autònoma de Barcelona

Vincent Iehlé[‡]

Université Paris-Dauphine

Nuria Rodriguez-Planas[§]

Universitat Autònoma de Barcelona

March 15, 2010

[Preliminary and incomplete]

*The first and third authors thank the *Barcelona GSE Research Network*, from the Government of Catalonia. Support from *Plan Nacional I+D+I* (SEJ2005-01481, SEJ2005-01690), the *Generalitat de Catalunya* (SGR2005-00626) are gratefully acknowledged.

[†]Corresponding author: Guillaume.Haeringer@uab.es.

[‡]vincent.iehle@dauphine.fr.

[§]Nuria.Rodriguez@uab.es.

Abstract

The entry-level French academic job market is organized by a central clearing house, where departments and candidates must submit their preferences over potential matches. We show that the algorithm used in France is equivalent Gale and Shapley's (1962, *American Math. Month*) Deferred Acceptance algorithm with candidates proposing. Using data from 1999 to 2007 of the market for mathematicians we observe that departments strategically divide the market into two sub-markets, one in which candidates' preferences are "knocked out" (i.e., their preferences are not needed to deduce the final match) and one in which the final outcome depends on candidates' preferences. The average quality of the matches turns out to be significantly affected when candidates' preferences are knocked out. In particular, those candidates publish significantly less in the years after having been hired. When competition is low between departments (i.e., many top-tier university are not hiring) low-tier departments manage to hire top-tier candidates, whereas for high levels of competition between departments knocking out candidates' preferences leads departments to be opt for a safety choice by targeting worse candidates.

KEY WORDS: French academic job market, centralized matching, market thickness, sorting, strategic behavior.

Classification of Journal of Economic Literature : C78, C62, J41.

1 Introduction

Over the last decades the market design literature, and, in particular, the matching literature have been able to understand better why some markets fail or succeed. Another lesson offered by the literature is that centralized clearinghouses can sometimes perform better than decentralized market institutions. As Roth (2008) explains, centralized markets may be able to correct many of the pitfalls we can observe in decentralized markets, like congestion.¹ Centralized clearinghouses, however, are not immune to market failures. McKinney, Niederle and Roth (2005) report for instance about the failure to organize the American gastroenterologist market, and identify the lack of market thickness as the most likely culprit. In this paper we document about another aspect that may debilitate (or affect) a matching market: constrained choice.

Before giving more details recall first that in a typical centralized matching market participants are asked to submit a preference list over potential partners and then a clearinghouse computes an assignment for each agent given the submitted preferences. Constrained choice in a centralized matching environment simply consists of limiting the number of items/potential partners that participants can put in the preferences they submit. This restriction may be exogenous (e.g., school choice procedures in New York City or College Admission in Spain) or it may be endogenous.² This latter case occurs for instance when it is costly (in money and/or time) for a firm to consider all potential partners. Note that when considering large markets constrained choice is more likely to be the norm than the exception. Constraining choices introduces an additional risk that market participants will have to face, that of being left unmatched. Using a controlled experiment Calsamiglia, Haeringer and Klijn (forthcoming) document about the distortion of market outcomes when agents become constrained. In particular, they observe that individuals tend to choose strategies that help them secure a minimal outcome, so as to minimize the risk of being left unmatched. The environment considered in this experimental study was that of a school choice market, where each student is endowed with a safe option, i.e., a school the student is certain to obtain if he puts it in the preference list he submits to the clearinghouse.³ The introduction of a safe option in a school choice market is feasible because in such markets students are the only active players. That is, schools do not choose their “pref-

¹See Roth and Xing (1997) for an account on the matter.

²See Haeringer and Klijn (2009) and Calsamiglia, Haeringer and Klijn (forthcoming) for a study of school choice matching when choice is constrained.

³In real-life settings this safe option is a school at which a student has the highest priority, which is usually a school that is located in the student’s home district. Depending on the algorithm used to match students the position of the district school (i.e., the “safe option”) in the submitted preference list may matter. For instance, in the so-called Boston mechanism the district school is a safe option only if put first in the ranking.

ferences” over students, and thus can be seen as mere objects to be consumed by students. It is thus very easy for the public authorities to tweak schools’ preferences over students such that for each student one of the schools is a safe option.⁴ When the market is such that both sides are active it may not be possible to secure a match for each participant, however. This is so even if we restrict to the short side of the market. Academic job markets are typical examples of such markets. The large number of participants and the costs associated to the recruiting of faculty (monetary or in time) significantly curbs the number of job candidates a department can consider. Also, academic departments do not have in general the possibility to secure a match (except that of not hiring).

The French academic job market turns out to be the perfect example of a centralized matching market where choice is constrained and both sides of the market are active players. In this market departments and candidates have to submit their preferences over potential hires to the Ministry of Higher Education, which computes the matching using these preferences. Until 2009 choice was constrained for the departments: they could not rank more than 5 candidates per position. Also, until 2009 French universities had no authority on the number of positions available on the market, which together with the fact that French departments typically lack faculty members (so as to cover departments’ teaching loads) makes unfilled positions very undesirable outcomes.

The main objective of the paper is to document how individuals tackle the constrain in a real-life environment, and to analyze the effects of choice restriction on market outcomes. Matching games are extremely complex games with, in general, a large number of Nash equilibria. The complexity stems from the fact that it is in general difficult to evaluate one’s prospects and strategy spaces in such games are very large (because a strategy in this context consists of an ordering over alternatives). Also, from the theoretical perspective, there is no clear prediction about what would be a likely outcome in a market like the French academic job market. We analyze in this paper data obtained from the French entry level market for mathematicians, between 1999 and 2007. Our data is obtained from a website ran by the community of mathematicians that keep tracks of all the decisions made by the departments (i.e., interview lists and rankings).⁵ We augmented our data set by including the final match (i.e., which candidate is hired by which university), and the publication records by departments and candidates —so

⁴To stress that schools are not agents and just “objects” the literature usually call schools’ preferences “priorities.”

⁵The website is similar in its spirit to the website used in the job market for economists <http://bluwiki.com/go/Econjobmarket>. There are two main differences, though. First, our data covers about 90–95% of the job openings, and second our data is extremely reliable —see Section 4.

as to have a measure of participants' (scientific) quality.

Presumably, a centralized clearinghouse needs preferences from both sides of the market to compute a matching. This implies that for both sides of the market there is a some uncertainty about the identity of the final partner (if any). The observation of departments' rankings in our data set significantly alters this intuition: for nearly half of the market we do not need candidates' preferences to compute the final matching. A typical case is when a candidate is ranked by only one department, and is ranked first. Clearly, this candidate ends up being hired by this department. In other words, for this department and this candidate there is no uncertainty (and thus no risk) with respect to the final outcome. The perfect predictability of (half of) the final matching thus cancels for half of the positions the risk of being unmatched. One crucial issue here is to understand which side is the one reducing the risks. Our analysis strongly suggests that this risk reduction is due to the departments' strategic behavior, and not from the candidates. That is, it is the departments and not the candidates that "knock out" candidates' preferences. We then face two sub-markets, one we may qualify as non-competitive (in the sense as only the labor demand is necessary to determine the final outcome), and one we may call competitive (where the supply and the demand are equally important).

A second stage of our analysis of the French job market consists of scrutinizing the effects of the division between two sub-markets. One may think that departments' preferences over candidates are mostly driven by candidates' qualities, such as the current publication lists when the candidate is on the market and their expected scientific productivity. Candidates' preferences are in contrast more prone to take into account non-academic aspects such as the cost of living, distance to the home city, etc. When candidates' preferences are knocked out, we may then expect that the matching is more "academic", i.e., more assortative. It is not. First, knocking out the supply in this labor market has an immediate effect on the quality of the match in the sense that most of the best candidates are matched in the competitive market. The picture is bleaker for departments, for good and average departments are as likely to go to either market. Although we observe some sort of assortative matching (in terms of the scientific productivity of departments and candidates), matchings in general are not assortative.⁶ Departments often hire candidates that are better or worse than the candidates that would correspond to their level. We find however that departments hiring on the non-competitive sub-market are more likely to mismatch. Further scrutiny of our data shows that the quality of the match in the non-competitive market depends on the degree of competitiveness across departments in a given year. Not all departments have a job openings each year. It follows that the ranking of a

⁶On possible source of the high level of noise in the match quality is perhaps the difficulty to assess the quality of a candidate at such an early stage of their career.

department in a given year (among the departments having at least one job opening) may not be equal to its “true” ranking. The higher is the difference between these two rankings, the less “competition” this department is facing. We find that the quality of the match in the non-competitive market is positively correlated with the absence of competitiveness. That is, the absence of good departments on the market gives more hiring opportunities for low-tier departments. The converse holds true: when competition is high departments on the non-competitive sub-market tend to hire candidates of lower quality. This is where the safety strategy takes place: A department seeking to secure a hire is more likely to succeed if it targets a candidate that is below its level.

The paper is organized as follows. We first review in Section 2 the basic setup of a matching model. In Section 3 we describe the French academic job market. A precise account of the data we obtained is given in Section 4. In Section 5 we document about the market divide between a sub-market where the outcome can be perfectly predicted by just knowing departments’ submitted preferences and another sub-market where both sides’ preferences are needed. In Section 6 we analyze the consequences of this divide on market outcomes.

2 Preliminaries

A typical matching problem consists of a (finite) set of workers and a (finite) set of firms, $W = \{w_1, \dots, w_m\}$ and $F = \{f_1, \dots, f_n\}$ respectively. Each worker $w \in W$ has a strict, complete, transitive, and asymmetric preference relation P_w over $F \cup \{w\}$. Similarly, each firm $f \in F$ has a strict, complete, transitive, and asymmetric preference relation P_f over $W \cup \{f\}$. For each player $v \in F \cup W$, let R_v denote the weak preference relation associated to P_v . The set of all possible preferences for a player $v \in W \cup F$ is denoted by \mathcal{P}_v , and an element P of

$$\mathcal{P} \stackrel{\text{def}}{=} \prod_{w \in W} \mathcal{P}_w \times \prod_{f \in F} \mathcal{P}_f$$

is called a *preference profile*. A *job market* is described by the triple

$$(W, F, P).$$

A *matching* μ is a function from $W \cup F$ into itself such that

- (i) For all $w \in W$, $\mu(w) \in F \cup \{w\}$;
- (ii) For all $f \in F$, $\mu(f) \in W \cup \{f\}$;
- (iii) For all $v \in W \cup F$, $\mu(\mu(v)) = v$.

Given a matching μ , we will refer to $\mu(v)$ as the *mate* of v under μ . Let \mathcal{M} denote the set of all matchings. Given a preference relation P_w of a worker w over $F \cup \{w\}$, we can extend P_w to a complete, transitive and reflexive preference relation R_w over the set of matchings \mathcal{M} in the following way (abusing notation); $\mu P_w \mu'$ if and only if $\mu(w) P_w \mu'(w)$ and $\mu R_w \mu'$ if and only if $\mu(w) R_w \mu'(w)$.⁷ Similarly, we can also extend firms' preferences to preferences over \mathcal{M} .

2.1 Stability and Deferred Acceptance Algorithm

For a firm f (resp., worker w), a worker w (resp., firm f) is *acceptable* if f (resp. w) prefers to be matched to w (resp. f) rather than remaining unmatched, i.e., $w P_f f$ (resp. $f P_w w$). For a player $v \in F \cup W$, $A(P_v)$ denotes the set of acceptable players with respect to the preference relation P_v .⁸

A matching μ is *individually rational* at a profile P if for each player $v \in F \cup W$ it holds that $\mu(v) R_v v$, i.e., $\mu(v) \in A(P_v) \cup \{v\}$. The set of individually rational matchings at P is denoted $IR(P)$. A matching μ is *blocked by a pair at P* (w, f) if $f P_w \mu(w)$ and $w P_f \mu(f)$. A matching μ that is individually rational and is not blocked by any pair is called *stable*. We denote by $S(P)$ the set of stable matchings at P .

A stable matching can be obtained using Gale and Shapley's (1962) *Deferred Acceptance* (DA) algorithm. Their algorithm, with workers making proposals to firms, works as follows:

Step 1: Each worker w proposes to his most preferred firm among the ones that are acceptable for him. If there is no such firm then the worker is matched to himself.

Each firm declines all but its most preferred worker among the workers who proposed to this firm and are acceptable to it (if any).

Step k , $k \geq 2$: Each worker who has been declined in the previous step proposes to his most preferred firm among the firms that have not yet declined him and are acceptable for him. If there is no such firm then the worker is matched to himself.

Each firm declines all but its most preferred worker among the workers who proposed to this firm and are acceptable to it and the worker it did not decline in the previous step (if any).

⁷Preferences over mates are assumed to be strict. However, since at two different matchings an agent can be matched to the same individual, and thus be indifferent between these two matchings (i.e., there are no externalities), strict preferences over mates naturally extend to weak preferences over matchings.

⁸Note that by convention we set $v \notin A(P_v)$, for all $v \in W \cup F$.

The algorithm stops when every worker is either matched to a firm or to himself. Given a preference profile P , we denote by $\mu_W(P)$ the matching obtained by the DA algorithm we just described and by $\mu_F(P)$ the matching that is obtained when we invert the roles of workers and firms in the algorithm, i.e., firms propose to workers and workers reject or accept workers' proposals.

Theorem 1 (Knuth, 1976) *The matching $\mu_F(P)$ (resp. $\mu_W(P)$) is the workers' (firms') least preferred matching among all stable matchings and $\mu_W(P)$ (resp. $\mu_F(P)$) is the workers' (firms') most preferred matching among all stable matchings.*

Henceforth, we shall refer to the matching $\mu_W(P)$ as the *worker-optimal matching* (W -optimal matching), and $\mu_F(P)$ as the *firm-optimal matching* (F -optimal matching).

3 The French job market

3.1 Calendar and main aspects

We consider in this paper only the junior positions, called *Maitre de Conférences*. They correspond to a lectureship in the U.K., or to an assistant professorship in the US. Junior positions in France are civil servant, tenured positions. Tenure is made official after one year of service, but it is extremely rare that a candidate is refused the tenure. Salary and teaching loads are, in principle, not negotiable. The salary is calculated given the rules of the ministry, possibly taking into account the years of service the candidate has done as a teaching or research assistant.⁹ The teaching load is 128 hours of lectures per year, independently of the university.¹⁰ Moving costs are not covered by the hiring department and access to research fund is often limited, if not available at all during the first years. All these features leave very little room for negotiation in the job market, and thus candidate's preferences over positions are exogenous to the job market outcome.

For candidates the job market begins in December or January. At this time candidates apply first for the right to participate to the job market. This consists of sending the Ph.D. dissertation, papers, and the résumé to a national committee, the *Conseil National des Universités* (CNU). One of the main purposes of this committee is to discard candidates that are too weak. A positive

⁹In 2008 the lowest starting salary was €1745 per month, before income taxes.

¹⁰It often occurs (in mathematics) that assistant professors teach section courses, in which case the teaching load is higher. The rate is 1 hour of lecture is equivalent to 1.5 hours of section. Full professors have the same teaching load, but rarely teach section courses. In few universities (e.g., Paris XI – Orsay) a reduction of the teaching load is offered to young scholars the first years.

evaluation is valid for four years, after which the candidate must ask for another assessment by this committee, should he or she want to participate again to the job market. Positions are sorted by fields, or *sections*. There are about 74 different sections, but for medical positions sections are often split into sub-commissions corresponding to sub-specialities, which brings the total number of commissions to 122. In this paper will shall consider the two sections related to mathematics, sections 25 and 26, denominated “mathematics” and “applied mathematics and mathematics applications,” respectively.¹¹ There are as many committees of the CNU as there are sections, and a candidate can ask for an evaluation by as many sections as he or she desires. Although a candidate having passed the evaluation by, say, section 64 (biochemistry and molecular biology), has the right to apply for a position open in, say, section 5 (economics), recruiting commissions usually only consider candidates that have been approved by the CNU committee of the section corresponding to the position. Needless to say, positions are often advertised for several sections.¹² It is important to note that a candidate must hold a Ph.D. degree when applying for the evaluation by the CNU. That is, the academic job market in France is only open to doctors, and not to “Ph.D. candidates.”

Until 2008, job openings are decided by the ministry of higher education, taking into account the voes expressed by each universities. Applications to a position can be sent as soon as it is published, around February or March.

To each open position is assigned a recruiting committee of about 15–20 people. This committee is mostly composed of professors from the faculty where the position is open, but also contains a few external members, usually professors from other universities. Upon receiving the candidates’ applications the recruiting committees decide first which candidate to interview. In the sequel we shall refer to the list of candidates to be interviewed as the *short list*. The recruiting committee then chooses a date at which candidates will be interviewed, usually in May. In most cases all the interviews (of about 20–30 minutes each) are made the same day, but sometimes the interviews last two or three days. At the end of the day, or the day after, the committee establishes a *rank list*, i.e., a ranking of the candidates, which will be submitted to the ministry. For each position, a committee can only rank five candidates. In some cases two or three positions at the same faculty are “linked,” meaning that the committee must give the same ranking for each of the linked position. In that case the maximal number of candidates

¹¹Section 25 refers to what mathematicians call “pure mathematics.”

¹²It could be for instance, there are some positions targeted for applied mathematicians doing economics (so the position is advertised in both the applied math and economic sections), or a position labelled “mathematical methods in quantic molecular chemistry” opened to sections 26 (applied mathematics) and 32 (organic, mineral and industrial chemistry).

that can be ranked is 5 multiplied by the number of linked positions. In the sequel we will refer to a candidate that appears in the ranking of a position as a *ranked candidate*.

It is important to note that recruiting committees do not coordinate when deciding the date (and often the time) at which they will interview the candidate. So a candidate may have an interview scheduled in Paris and Marseilles (900 km away) the same day, at the same time. Some committees are flexible and allow candidates to tune their schedules but lack of flexibility seems to be more the norm than the exception. Transportation costs incurred during the interview phase (and lodging, if needed) accrue to the candidates.

Once the interview season is over, candidates are asked to submit (usually in June) their preferences to the web site of the ministry of higher education. A candidate can only submit a preference ordering over the positions for which he or she is ranked. About 10 days later an assignment is published, and candidates have about a week to confirm whether they accept the position. A position that is refused is lost.

Remark 1 The ministry organizes a second job market in September. In some cases, the positions that were not filled in the first job market are proposed again, but this is not necessarily the case. This session is very small and will not be considered in this paper.

Remark 2 Parallel to the university job market there is another job market involving research institutions (CNRS, INRA, INSERM, INRIA, etc.). These positions (also tenured) are research only positions and are usually preferred by the candidates. These institutions usually announce which candidate is offered a job in April, i.e., before the interviews. However, candidates have little control over the institution where they will be assigned. The assignment is usually announced in June, when candidates must submit their choice lists for the assistant professor positions. That is, a candidate hired by the CNRS (in April) does not know before June *where* (e.g., Paris, Marseilles or Toulouse) he will be assigned.

Remark 3 There are each year between 80 and 130 positions open in sections 25 and 26. This high number of positions is rather unusual. For instance, in economics there are on average about 40–50 positions a year. This high number is partly explained by the fact that mathematics departments often provide the mathematics courses taught in other disciplines. Each year, there are on average between 400 and 500 candidates having passed the evaluation of the CNU of sections 25 or 26,¹³ so this gives an average of 5 candidates per each opening. On average about 30–40% of the candidates do not pass the CNU evaluation. The main reasons of reject are misfit

¹³There are slightly more candidates qualified for section 26 than for section 25. The former often includes candidates from other disciplines (economics, physics, chemistry, computer science, etc.)

(the field of specialization of the candidate does not match the field of the section), incomplete application and weakness of the scientific production (about 5% of the applicants).

3.2 The matching algorithm

The ministry of higher education uses candidates' submitted preferences and the rank lists to assign candidates to positions. The procedure is called *ANTARES*. Participants to the market, whether candidates or recruiting committees have in fact very little information about the exact procedure used by the ministry.

When logging on the website a candidate faces the following instructions:¹⁴

- Case 1:** *You are ranked only once and ranked first.* You can either accept or reject the position.
- Case 2:** *You are ranked several times and always first.* Being ranked first each time you can obtain any of these positions, so you do not need to state a preference over all positions.
- Case 3:** *You have been ranked several times, and among these positions you are ranked at least once first.* You must submit a preference ordering. If the first choice is a position for which you are ranked first this is enough. Otherwise, you will be awarded the position only if the candidates ranked before you declined the position. In case your first choice(s) is(are) not satisfied, the procedure guarantees you to be assigned to one of the position for which you are ranked first, possibly by putting this position as the last in your preferences. Not stating a preference over a position will be understood as a rejection of this position.
- Case 4:** *You have been ranked several times, but never first.* You must submit a preference ordering. You will be awarded the position only if the candidates ranked before you declined the position.

One immediately sees that these informations are not enough to infer the matching algorithm. However, a more careful reading of the instructions allow us to deduce the following properties that any matching produced by *ANTARES* should satisfy:

¹⁴The instructions contain in fact six cases. Case 5 concerns those candidates applying for an assistant professor position and a full professor position. In that case candidates are invited to submit two preference lists, one for the junior and one for the senior positions. Next, the candidate must say which type of position he prefers, i.e., whether junior or senior. So, candidates can only submit lexicographic preferences and not inter twinned preferences, e.g., professor at university *X* (1st choice), assistant professor at university *Y* (2nd choice), professor at university *Z* (3rd choice). We have not seen in the data the case of a candidate having the possibility to choose between junior and senior positions. The sixth case is for those candidates that are not ranked anywhere. For them there is nothing to do.

- (i) A candidate cannot be matched to a position if he or she rejected this position (so, the position is declared as unacceptable).
- (ii) A candidate is matched to a position only if the candidates ranked above him or her have denied the position.

In words, the ANTARES procedure produces a stable matching. What is more problematic is that any stable matching satisfies the ANTARES requirements.

A few year ago Balinski (2001) criticized the matching procedure employed by the ministry, understanding that the matching computed by ANTARES would be the faculty-optimal matching. The ministry replied to invalidate Balinski's hypothesis and provided a clear description of the algorithm that is used. The algorithm is a combination of three algorithms, MOD1, MOD2 and MOD3. We now give a brief description of each of these algorithms.

MOD1: Finding best-best matches.

Round 1: Match all the candidates who are ranked first at their first choice to their first choice.

Round k , $k \geq 2$: For each candidate and department that has not been matched at a previous round, delete from their choice lists and rank lists the candidates and departments that have been matched in the previous rounds. If a candidate's choice list does not contain any department, match this candidate to himself or herself. If a department's rank list does not contain any candidate match this department to itself. Finally, match all the candidates who are ranked the highest (in the reduced rank lists) at their most preferred choice (in the reduced choice lists) to this position.

End: The algorithm stops when either all candidates and departments are matched or when there is no pair of candidate-departments who are mutually their first choice.

MOD2: Cleaning the choice and rank lists.¹⁵ This algorithm is activated only if MOD1 is not able to find a match for all candidates (assigned to a position or unassigned) and all positions (assigned to a candidate or the position is not filled). This algorithm consists of deleting the choices and ranks that will never be expressed. Consider for instance a candidate ranked only at some department, and ranked, say, 3rd, and the corresponding position is the candidate's first choice. Since the matching that will be eventually produced must be stable, it cannot be that the candidate ranked 4th or 5th are awarded the position, for otherwise the resulting matching would not be stable. In this case MOD2 deletes from

¹⁵A formal presentation of MOD2 is offered in the Appendix.

the rank list the candidates ranked 4th and 5th, and for these candidates the position is deleted from their choice lists.

Once MOD2 is completed ANTARES re-run MOD1. It may be that MOD2 will be activated a second time, then MOD1 a third time, then MOD2, and so on.

MOD3: Computing the candidate-optimal matching. This step is only activated when the combination of MOD1 and MOD2 (possibly repeated several times) does not produce a matching.

Interestingly enough, MOD3 is almost never activated.¹⁶ One issue that arises is what are the properties of the matching when the combination of three algorithms is used. It turns out that the ANTARES algorithm is equivalent to the Deferred Acceptance algorithm with candidates proposing. The intuition is rather simple. The algorithm MOD1 computes matches that are part of any stable matching.¹⁷ As for MOD2, it turns out that it does not affect the set of stable matchings given the choice lists and the rank lists. So, the main “difficulty” consists of showing that MOD1+MOD3 is equivalent to running MOD3 only. We then have the following result, the proof of which can be found in the Appendix.

Theorem 2 *For any profile of choice lists and rank lists, the ANTARES algorithm is equivalent to the Deferred Acceptance algorithm with candidates proposing.*

4 The data

4.1 Opération Postes

In 1998 a small group of young mathematicians set up a web site, *Opération Postes*, inviting recruiting committees to announce the short lists as well as the rank lists as soon as these would be decided.¹⁸ The community of mathematicians was very responsive and the web site quickly became a central tool in the job market. The web site is now supported by the main professional mathematical societies in France.¹⁹

¹⁶Considering all disciplines, MOD3 was activated less than 20 times in the last 20 years —Private communication with the person in charge of the maintenance of the ANTARES computing center.

¹⁷When MOD1 can compute the entire match (i.e., without invoking MOD2 nor MOD3) the choice lists and the rank lists have a property similar (but weaker) than Banerjee, Konishi and Sönmez’ *top coalition property*. It is well known that in this case the set of core-stable coalition structure is a singleton.

¹⁸<http://postes.smai.emath.fr/>.

¹⁹These associations are the French Mathematical Society (SMF), the Society for Applied and Industrial Mathematics (SMAI), which hosts the web site, the French Society of Statistics (SFdS),

Each year, the list of openings published by the ministry is posted on the web site and the presidents of the recruiting committees are invited to post the short list and the rank list as soon as these are decided. On average, about 90–95% of the openings’ short lists and rank lists are available.

The website also collects some additional information that may be of interest for the candidates, such as the position profile (i.e., whether the position is targeted towards a specific sub-field), the research center where the hired candidate will be assigned.²⁰ The web site also invites the recruiting committees to submit the date when the short list will be decided, the date of the interview, and the date when the rank list will be decided. This information is more sparse, for many committee presidents only submit the short list and the rank list. Until 2002, candidates who would decide to take a job for which they are ranked first were invited to announce it. In general, about 50% of the candidates ranked first made such announcements. We were also able to retrieve, for most years, the date at which the recruiting committees posted their lists. More generally, we do not have the choice lists of the candidates. Such lists are confidential and the ministry of higher education denied us the access to this data. So our data is only one-sided. However, we shall show that this information is not necessary to clear about half of the market.

The data of *Opération Postes* is public, but unfortunately not in a format that makes it immediately usable for any analysis. There are many misspellings, and in we often found confusions between the married and maiden names of some female candidates.

Table 1 provides an overview of the data available to us. The column “Missing” displays the number of openings for which we do not have any data. About half (per year) of these missing job openings are openings for which the recruiting committee did not publish the ranking list, the other half are positions that are not available for most of the candidates.²¹ The other columns give the number of positions available each year per section for which we have at least the rank list, and the number of candidates being shortlisted by for at least one position.²²

²⁰Research in universities is made in “laboratories” or research centers, the departments (called *Facultés*) are essentially designed to organize the teaching.

²¹These positions are either reserved to holders of a “assistantship” position (an old status that does no longer exist, akin to a “lecturer” in the US, or to individuals who are already assistant professor in another university and ask for a transfer.

²²Recall that we do not have the data for *all* available positions, but only 90–95% of them.

Year	Numb. of positions				Numb. of candidates		
	Missing	Sec. 25	Sec. 26	Sec 25& 26	Total	shortlisted	ranked
1999	12	35	65	20	120	439	253
2000	20	20	52	17	89	367	202
2001	22	25	38	23	86	339	201
2002	29	21	31	18	70	335	168
2003	22	27	52	18	97	383	214
2004	7	24	38	11	73	361	191
2005	10	22	41	26	89	399	210
2006	14	27	70	25	123	445	265
2007	20	35	61	19	115	471	267

Table 1: Number of positions and candidates per year

4.2 Additional data

We completed the database in several directions. First, for each candidate that is ranked at least once we searched on the Internet where the candidate made the Ph.D. and the year of defense.²³ Second, we identified for each position who was awarded the job. For some years this information is offered by the *Opération Postes*. Each year some candidates have an “outside option,” which is activated. An outside option consists of holding an offer from another institution, such as a research position or a foreign academic institution. It is often the case that such candidates having such offers do not even participate to the interview phase.²⁴ However, some candidates can and do wait to the end of the interview to decide whether activating their outside option. That is, these candidates appear in the rank lists (sometimes in the first position), but are not awarded the job. From the data we have it is difficult to know whether the market identify candidates with an outside option. The fact that we observe such candidates may be due either to a lack of information from part of the departments that rank them or because they just take their chance.

²³The main sources of this information were the *Mathematics Genealogy Project*, (<http://genealogy.math.ndsu.nodak.edu/>) and the French Ph.D. dissertation database in mathematics (<http://math-doc.ujf-grenoble.fr/Theses/>). In some cases we had to look up the candidate’s papers or web page to retrieve the information about the Ph.D.

²⁴This is sometimes due to the fact that the other offers are exploding, i.e., the candidate has a deadline to answer, which may be before the ANTARES assignment is produced.

4.3 Publications records

We also retrieved some information related to the scientific production of the actors of the market. We only considered the database *Mathscinet*, a database maintained by the *American Mathematical Society*.²⁵ This database covers a wide range of journals, such as computer science, physics, economics, psychology, or medicine.²⁶ Each publication was matched to the corresponding *impact factor* of the year corresponding to the publication year, which was obtained from the *Journal of Citation Reports* (JCR). To obtain impact factors we used the on-line version of the JCR.

The information relative to the publications is essential to us, for it enables us to obtain a ranking of the mathematics departments in France. Quite surprisingly such a ranking does not exist in France.²⁷ Since we wanted to avoid fetching publication records by non-math departments but indexed in Mathscinet we had to use the precise affiliation of the papers published. However, the lack of norms for writing the affiliation and the division of departments in smaller entities (the laboratories) considerably increased the number of institutions we had to look up in Mathscinet. In our case this amounts to consider more than 700 different affiliation codes.²⁸

We analyzed the publication records made by each of the institutions that have hired at least once between 1999 and 2007, which amount to 97 different institutions. There are 10 recruiting institutions for which we did not retrieve the publication records, the so-called *IUFM* (Institute for the formation of school masters). The teaching load (as well as the salary) is the same as for the “regular” university departments, but since classes are small lecture hours are not regarded as full lectures but rather section classes. That is, the teaching load in these institutes is in fact of 192 hours per annum (i.e., about 10 hours per week). Those institutions are known to be teaching oriented, with little research activity.²⁹ Three other institutions for

²⁵<http://www.ams.org/mathscinet/>. The other main database in Mathematics is *Zentralblatt MATH* (<http://www.zentralblatt-math.org/zmath/en/>).

²⁶Non-math journals are covered only if they are “mathematically oriented.” Journals like *Econometrica*, *Games and Economic Behavior* are covered but not *the American Economic Review* or the *Quarterly Journal of Economics*.

²⁷Some mathematicians explained us that a ranking is not seen as relevant, for saying that, say, University *X* ranks above University *Y* would amount to say that differential equations rank above geometry, the fields in which *X* and *Y* are respectively known to be strong.

²⁸The average number of different affiliation codes is between 5 and 10. For some universities (e.g., Paris VI or Paris VII) we identified nearly 50 different affiliation codes.

²⁹We did check the publication records of some of these institutions and they are significantly lower than any of the other institution we considered. So putting the IUFM’s last in our ranking is without loss of generality. The main reason being that mathematicians in IUFM’s doing research usually have another affiliation to a research laboratory in another institution or department.

which we do not have the publication records are the *Université de Nouvelle Calédonie* (near New Zealand), the *Université des Antilles* (Caribbean) and the *Université de La Réunion* (near Madagascar). The database *Mathscinet* does not provide any institution code for these three universities.³⁰ However, the geographical distance between these universities and continental France make those institutions less attractive for continental candidates. The candidates hired at those institutions are usually either local candidates or continental candidates with no other hiring option.

We considered only the papers published in journals between 1997 and 2007 (to match the impact factor given by the JCR). The total number of publication is steadily increasing, from about 2015 in 1997 to nearly 3000 in 2007 (there is a drop in 2008, with only 2479 publications). Our ranking spans 27,966 publications.

Impact factors are delicate to manipulate for two reasons. First, there is not a great variance in the impact factors of most mathematical journals. For instance, *Annals of Mathematics*, which is considered as best journal in pure mathematics has an impact factor of 1.5 on average. To correct this problem we considered the square values of the impact factors. A second issue is the presence of multidisciplinary mathematic journals, with impact factors that are in general much higher than those found in math-only journals. For instance, we find in our data set *Econometrica* or *Biostatistics* with average impact factors of 2.2 and 4.5, respectively. To correct for this sub-field effect we had to separate candidates and departments by subfields (e.g., statistics, geometry, analysis, etc.).

5 Market clearing

The first issue we raise is whether we have enough information to clear the market. There are two polar cases in a matching market, one in which we do not need to know the choices of any of the candidates and one in which we need to know all the choices. The former case arises for instance when the rank lists of any two departments have an empty intersection. In this case all the candidates that are ranked first are assigned the corresponding position (provided these positions are seen as acceptable by the candidates), i.e., departments avoid vying with each other. Not needing to know candidates' choice lists to clear the market suggest that there is a high degree of coordination between departments.

Our data shows that the set of candidates can be split into two different sub-samples, de-

³⁰The math departments of these universities are not significantly smaller than those in continental France (except for New Caledonia, with less than 10 members). A few searches using the names of these department members show decent publication records.

pending on whether department compete to hire them.

Recall that we do not observe candidates' choices submitted to the ANTARES web site. However, a careful look at the rank lists allows us to deduce some matches. This is the case for instance when a candidate is ranked first at some department and is ranked only once. In this case it is natural to conjecture that this candidate would take the job. Whenever we would deduce that a candidate would be assigned a position at some department we checked whether it was indeed the case. All assignments deduced by our strategy were confirmed by the data, which suggests that candidates see all positions as acceptable. Once a position is awarded to a candidate ranked first we deleted for the candidates ranked 2nd, 3rd, etc. that position from the list of positions at which they are ranked. We can also proceed in a similar fashion with any candidate that is ranked only once although being ranked, say, 3rd. In this case the candidates being ranked 4th and 5th are not eligible for the position, and thus we can eliminate the position considered in these two candidates' list of departments that ranked them. Such elimination of position in candidates' list of the departments that ranks them can in turn generate new situations in which a candidate is only ranked once.³¹ Iterating this procedure several times we eventually end up with a set of position already filled.³² Table 2 shows the number of positions filled this way.

year	pos. filled	% of all positions
1999	39	32.5
2000	45	50.5
2001	33	38.4
2002	40	57.1
2003	49	50.5
2004	36	49.3
2005	37	41.5
2006	41	33.6
2007	52	45.2

Table 2: Positions filled without needing candidates' choices

The first question that arises is how this separation in two sub-markets operates. One possibility is that departments somehow coordinate (perhaps implicitly) and opt for not ranking the

³¹Notice that this procedure somehow mimics (albeit with fewer information) the algorithm MOD1.

³²Usually one round of this procedure is enough to assign all positions that can be assigned without knowing candidates' preferences.

same candidates. This is perfectly feasible since in the job market under consideration departments happen to observe in real time the decisions of all other departments. That is, it is not difficult for a department to pick candidates that are not ranked by other departments. Another possibility is that there is a self-selection effect. A candidate without good prospects may indeed want to withdraw from the job market once he has secured a position at some department. This behavior can be explained by the fact that in the French job market transportation and lodging costs accrue to the candidates.

The best way to assess whether candidates are sorted out by departments is to see whether candidates in the non-competitive market do show up to the interviews scheduled *after* the interview for the position they eventually obtained (in which case the market divide would rather be due to the departments). The problem with this strategy is that we do not observe interview attendances. Not appearing on the rank list for some position in spite of being scheduled for an interview may be due either to the fact that the department preferred to rank other candidates, or to the fact that the candidate did not come to the interview. Note, however, that a candidate appearing in one ex-post ranking is a sufficient information to deduce that the candidate did not self-select.³³ It is important to be careful when looking at the number of times a candidate is ex-post ranked, since candidates in the two sub-markets do not have the same number of scheduled interviews. Candidates in the competitive sub-market are scheduled on average for 7.99 interviews, which easily compares with the other candidates' 6.46 interviews on average.³⁴

To circumvent the above mentioned issues we compared the number of ex-post scheduled interviews with the number of ex-post rankings. Self-selection cannot be discarded if we observe a high difference between the average ratio of these two numbers across the two sub-markets. Table 3 displays the statistics obtained when considering the ratios ($\#$ of ex-post rankings)/($\#$ of ex-post interviews) across the two sub-markets. We consider only the candidates that have at least one scheduled ex-post interview.

	Observa.	Average	Std. Dev.
Competitive sub-market	357	.481	.2564
Non-competitive sub-market	228	.456	.2478
<i>t</i> -stat/ <i>F</i> -stat (<i>p</i> -value)		1.1716 (.1209)	1.071 (.2881)

Table 3: Presence at ex-post interviews

³³In the sequel the term “ex-post” (interview, ranking, etc) will refer to the interviews that take place after the interview for the position a candidate eventually obtained.

³⁴The difference is highly significant: *t*-stat=4.16, *p*-value=.000.

The results displayed in Table 3 strongly suggest that the sorting between the two sub-markets is mostly due to departments’ choices and not a candidates’ self-selection. Another piece of evidence to sustain this thesis is obtained when analyzing the structure of departments’ rank lists. To start with, note that a department’s decision to hire on the non-competitive sub-market can be understood as a decision to hire a “sure-candidate,” as opposed to the other market where departments have less discretion over the candidate that will be eventually hired. In this job market, a sure hire is obtained when a department ranks candidates that are not ranked by any other department. In this context the number such candidates that appear in a department’s ranking is a good signal of the department’s precautionary behavior. To get a sense of this behavior we considered for each departments’ rankings the ratio of candidates that are not hired by any other departments. These candidates can be seen as candidate that count for nothing, i.e., they only appear in the rank lists for filling purposes. Table 4 summarizes the statistics obtained across both sub-markets.

	Observa.	Average	Std. Dev.
Competitive sub-market	478	40.45	20.49
Non-competitive sub-market	339	49.25	26.25
<i>t</i> -stat (<i>p</i> -value)		5.3757 (.000)	.609 (.000)

Table 4: Percentages of candidates counting for nothing in departments rankings

6 Sorting

The fact that sorting between competitive and non-competitive sub-markets is mainly due to departments’ strategic choice will have profound implications when analyzing the market outcomes. Departments’ preferences over candidates are presumably dictated by candidates’ qualities, in contrast with candidates’ preferences that are likely to be influenced by other factors, such as the distance to the university of origin, the cost of living or the attractiveness of the region where departments are. One would then expect that there is more sorting in the non-competitive sub-market, where the outcome solely depends on departments’ choices, than in the other sub-market.

To analyze the effects of the market divide we described in Section 5 we first compare the pools of candidates that are matched in each of these sub-markets. To this end, we first sum for each candidate hired in year y the squares of the impact factors his publications in years $y + 1$,

y , $y - 1$ and $y - 2$.³⁵ Taking the squares of the impact factor is to exacerbates the differences between journals in mathematics, a discipline with a very low variance in the impact factors. We do not use a strict ranking of candidates in terms of their publications records. Instead, we consider the quartiles, with the first quartile containing the candidates with no or few (and low quality) publications, and the fourth quartile containing the best candidates.

Table 5 displays the distribution of candidates' quality for the two sub-markets. For top half of the market (3rd and 4th quartiles) the competitive sub-market outperforms the non-competitive submerged, for the latter matches not only more of these candidates but also in proportion.

	Observ.	1st quartile	2nd quartile	3rd quartile	4th quartile
competitive	478	22	24.7	27.2	26.2
non-competitive	339	31.6	27.1	22.4	18.9
<i>t</i> -stat (<i>p</i> -value)		3.09 (.001)	.79 (.215)	1.55 (.0608)	2.43 (.0076)

Table 5: Distribution of candidates' quality

The distribution of candidates' quality takes a different picture when we relate it with the degree of competitiveness between departments. The degree of competitiveness here is measured by the difference between the "true" ranking of a department and its ranking in a given year, if present on the market. Since not all departments hire every year, a part of the difficult to hire for a given department can be captured by measure the absence (or presence) of department that are better. The average degrees of competitiveness between both sub-markets are substantially different. In the non-competitive market a department increases on average its rank by 12.85 positions, which is significantly higher than the 8.58 increase we observe for the competitive sub-market.³⁶

The analysis becomes however more instructive if we distinguish between the departments that manage to hire a better candidate from those that do not. The concept of "better" or "worse" candidate we use is simply by comparing the quartiles of the department and that of the candidate. That is, a candidate in the k -quartile hired by a department in the k' -th quartile is a better candidate for this department if $k' < k$. Conversely, he is a worse candidate if $k' > k$.

Table 6 helps understanding how departments' relative success (i.e., hiring a better or worse candidate) is related to the degree of competitiveness in the market. The second and third

³⁵Publication delays (review process and backlogs) in mathematics are similar to those present in economics. Since the market takes place in May papers published in year $y + 1$ are almost always already written in May, year y .

³⁶The *t*-stat and *p*-value are 3.0704 and .0011, respectively.

columns of Table 6 give the average increase in ranks of departments that hire a better candidate (“Yes”) and of departments that do not hire a better candidate (“No”). It is clear that lack of competition is closely associated to the possibility to hire a better candidate. Columns 4 and 5 provide similar statistics but considering departments that hire a worse candidate.

	Hiring better candidate		Hiring worse candidate	
	Non-comp. market	Comp. market	Non-comp. market	Comp. market
No	5.09	4.27	16.84	11.22
Yes	25.83	15.74	9.06	5.25
<i>t</i> -stat (<i>p</i> -value)	7.7735 (.000)	6.2864 (.000)	5.3757 (.0033)	3.5043 (.0003)

Table 6: Competitiveness depending on departments’ relative success

Departments in the non-competitive market hiring a better candidate experience a significant increase of their relative ranking than departments in the competitive market hiring a better candidate (25.83 *vs.* 15.74). This difference is not significant for departments that do not hire a better candidate.³⁷ Table 6 suggest then that departments in the non-competitive market experience a greater inertia, in the sense that they need a greater change in their relative rankings than departments in the competitive market to hire a better or a worse candidate. This reading of Table 6 can be misleading, for it does not take into account that departments in the non-competitive market do experience in general a greater change in their relative ranking (12.85 and 8.59 for the non-competitive and competitive markets, respectively). If we take this into account we find that departments in the non-competitive sub-market are more likely to hire a better or a worse candidate than departments in the competitive sub-market. In other words, the competitive sub-market is more assortative than the non-competitive sub-market. The difference is no significance when considering the recruitment of better candidates, but it is significant for worse candidates. Table 7 presents these findings.³⁸

	Hiring a worse cand.	Hiring a better cand	Same quality
Competitive sub-market	28.33	54.46	25.7
Non-competitive sub-market	31.82	64.25	18
<i>t</i> -stat (<i>p</i> -value)	2.4225 (.0078)	.8573 (.1958)	2.6166 (.0045)

Table 7: Percentages of firms hiring a better or worse candidate.

³⁷Comparing 5.09 and 4.27 we obtain a *t*-stat equal to .9711 (*p*-value=0.1921).

³⁸To compute the percentage of departments that hire a better or a worse candidate we did not consider the departments that are in the fourth and first quartile, respectively.

The last columns of Table 7 shows that assortative matching (i.e., department in the k -th quartile hiring a candidate in the k -th quartile) is much lower than one would expect in this market, which is known to be fierce. Nevertheless, we do observe, as we commented above, that the competitive sub-market is more assortative than the non-competitive sub-market.

An important question that is left answered is how departments select the candidates that will be hired on the competitive or non-competitive sub-market. An important part of the answer lies in the candidates' scientific productivity after having been hired. To this end we consider for a candidate hired in year y the publications in year $y + 2$, $y + 3$ and $y + 4$. Taking into account the fact that our measure of candidate's productivity at the time they are hired takes into account 4 years of (potential) publications, we used the following index to measure the progress in terms of scientific production

$$progress = \frac{4\text{pub. ex-post} - 3\text{pub. ex-ante}}{3\text{pub. ex-ante}} \text{ if pub. ex-ante} > 0.$$

	3rd-4th quartile dept.	1st-2nd quartile dept.	ex-post publications
Competitive sub-market	4.17	4.33	1.95
Non-competitive sub-market	2.24	12.73	1.5705
t -stat (p -value)	1.448 (.0747)	1.74 (.0415)	1.295 (.0978)

Table 8: Candidates' progress in scientific production

Table 8 shows the differences between the competitive and non-competitive sub-markets. We first observe (4th column) that candidates hired on the competitive sub-market publish more than those in the non-competitive sub-market. This suggests that departments vie with each other for the candidates with a higher potential, or at least those with a research agenda that will lead to publications. This interpretation is confirmed by the observation that top departments (3rd and 4th quartile) do indeed perform better at hiring the candidates with a higher potential in the competitive sub-market (2nd column). What is perhaps more surprising is the fact that low-tier departments are able to find extremely promising candidates in the non-competitive market. One explanation we found when communicating with the mathematician community is that low-tier departments do spend more time examining candidates that are apparently candidates with a low profile during the job market.

7 Discussion

Our analysis of the strategies and the outcomes of the French job market for mathematicians also contributes to the matching literature in the sense that it sheds some light on the type of behavior

we can expect to observe in a centralized matching market. Since the algorithm used by the French ministry is equivalent to the Deferred Acceptance algorithm with candidates' proposing, we know from the theoretical literature that it is a dominant strategy for the candidates to reveal their true preferences —Dubins and Freedman (1981) and Roth (1982). However, it is not strategy-proof for the departments. Roth (1984) and Gale and Sotomayor (1985) showed that the set of Nash equilibria can be quite large, which leaves us with virtually no sharp prediction about the type of strategies we may observe in a real-life settings. Analyzing the different rankings for each year, we find very few contradictions, i.e., the relative rankings of the candidates are almost always rationalizable by a ranking according to candidates' publication records. In other words, we may assume without too much risk that departments' preferences over candidates coincide and are driven by candidates' qualities. In the market we analyzed departments cannot rank many candidates and not hiring is very costly. So, the issue is not rather to understand how departments would manipulate their preference but rather what are the candidates they target when on the market. We first observe that departments manage easily to avoid their positions unfilled.³⁹ Given that the web site *Opération Postes* allow departments to observe the decisions taken by their competitors, this fact is hardly surprising.

The fact that departments are constrained in their choices resembles what is known in the matching literature as *truncation strategies*, initially studied by Roth and Rothblum (1999). Truncation strategies are useful in Deferred Acceptance games to reverse the outcome for the side receiving the proposals. In this case it is well known that by truncating their preferences the agents receiving the offers (the departments in our case) by declaring as unacceptable any candidate less preferred than the candidate they obtain at the department-optimal matching can manage to be matched to their department-optimal candidate. One could then expect that in the French job market, the constrain of ranking at most 5 candidates implies that departments are would yield the departments to adopt a strategy that would yield the department-optimal matching. Our results suggest that this is not the case, for most departments end up being matched to a worse candidate.

³⁹We observed only two positions not filled over the 9 years we analyzed.

A Appendix: Proofs

We first give a formal presentation of the algorithm MOD2, and then show that it does not affect the set of stable matchings.

Let $P_C = (P_c)_{c \in C}$ and $P_F = (P_f)_{f \in F}$ be the profiles of choice lists and rank lists, respectively. Let $P = (P_C, P_F)$ and let $S(P, C, F)$ denote the set of stable matching under the profile P when the set of candidates is C and that of departments is F . Let c_0 be a candidate and f_0 a department such that c_0 is ranked by f_0 and f_0 is c_0 's first choice, i.e., $f_0 P_{c_0} f$ for each $f \neq f_0$ such that $f P_{c_0} c_0$. Let \tilde{c} be a candidate such that $c_0 P_{f_0} \tilde{c} P_{f_0} f_0$. MOD2 consists of defining a new profile P' satisfying the following:

- For each $c \neq \tilde{c}$, $P'_c = P_c$ and for each department $f \neq f_0$, $P'_f = P_f$.
- For any $v, v' \in (C \cup \{f_0\}) \setminus \{\tilde{c}\}$, $v P_{f_0} v'$ if, and only if, $v P'_{f_0} v'$, and $f_0 P'_{f_0} \tilde{c}$.
- For any $v, v' \in (F \cup \{\tilde{c}\}) \setminus \{f_0\}$, $v P_{\tilde{c}} v'$ if, and only if, $v P'_{\tilde{c}} v'$, and $\tilde{c} P'_{\tilde{c}} f_0$.

This transformation is repeated as many times as we can find a pair of candidates and a department in the same situations as c_0 , \tilde{c} and f_0 .

Proposition 1 *Let P be a profile of preferences and P' the profile obtained by MOD2 with input P . $S(P, C, F) = S(P', C, F)$.*

Proof Clearly, $\mu \in S(P)$ implies $\mu \in S(P')$. To see this, notice that if (c, f) blocks μ under P' then (c, f) also blocks μ under P . Let $\mu \in S(P')$ and suppose $\mu \notin S(P)$. Since $\mu \in S(P')$, $\mu \in IR(P')$. By construction, we also have $\mu \in IR(P)$. So, $\mu \notin S(P)$ implies that there exists a pair (c, f) that blocks μ under P , i.e.,

$$c P_f \mu(f) \quad \text{and} \quad f P_c \mu(c). \quad (1)$$

Suppose $f \neq f_0$. So $P_f = P'_f$. If $c \neq \tilde{c}$ then $P_c = P'_c$. It follows that (1) is tantamount to $c P'_f \mu(f)$ and $f P'_c \mu(c)$. We then have $\mu \notin S(P')$, a contradiction. So, $c = \tilde{c}$. Since $\mu \in S(P')$ and $\tilde{c} P'_{\tilde{c}} f_0$, $\mu(c) \neq f_0$. Hence, $f P'_{\tilde{c}} \mu(\tilde{c})$ is equivalent to $f P_{\tilde{c}} \mu(\tilde{c})$, and thus the pair (\tilde{c}, f) block μ with respect to P' , which contradicts $\mu \in S(P')$. So, $f = f_0$. Since $\mu \in S(P')$, $\mu \in IR(P')$ and thus $\mu(f) R_f c_0$. So (1) implies $c P_f c_0$ and thus $c \neq \tilde{c}$. We then have $P_c = P'_c$ and $c P'_f \mu(f)$, which implies that (c, f) blocks μ under the profile P' . This contradicts $\mu \in S(P')$. ■

Proof of Theorem 2

Step 1. Let P be a preference profile and let (c, f) be a pair matched by MOD1. Then for any matching $\mu \in S(P)$, $\mu(c) = f$.

Let (c_1, f_1) be a pair matched in the first round of MOD1, i.e., c_1 's first choice is f_1 and c_1 is ranked first by f_1 . Clearly, if $\mu \in S(P)$ then $\mu(c_1) = f_1$. Let (c_2, v) be a pair matched in the second round of MOD1.⁴⁰ Let $\mu \in S(P)$ such that $\mu(c_2) = v' \neq v$. We claim that $vP_{c_2}v'$. To see this, suppose that $v'P_{c_2}v$. Since c_2 is matched in round 2 of MOD1 to v it follows that v' is matched in round 1. So there exist $\mu, \mu' \in S(P, C, F)$ such that $\mu(v') \neq \mu'(v')$, a contradiction. So $vP_{c_2}v'$. Let $v'' = \mu(v)$. Using a similar argument we deduce that c_2P_vv'' . So the pair (c_2, v) block μ , which contradicts $\mu \in S(P)$. It suffices to repeat the argument with pairs matched at the next steps (if any) to obtain the desired result.

Step 2. By Proposition 1, we can consider without loss of generality that ANTARES is a combination of MOD1 and MOD3. Also, since MOD3 computes the candidate-optimal matching we can assume that MOD3 is the DA algorithm with candidates proposing. Let μ be the matching obtained by ANTARES and μ_C the candidate-optimal matching. Consider MOD1+MOD3 and let C' and F' be the set of candidates and departments matched during the MOD3 phase of ANTARES. If either $C' = \emptyset$ or $F' = \emptyset$ then MOD1 computes a matching, which is necessarily stable by Step 1, and thus $\mu = \mu_C$.

Suppose then that $C' \neq \emptyset \neq F'$. We denote by $\mu|_{C' \cup F'}$ the matching μ restricted to the agents in $C' \cup F'$, and by P' the profile P restricted to the agents in $C' \cup F'$. We now show that $\mu \in S(P, C, F)$. Notice first that $\mu|_{C' \cup F'} \in S(P', C', F')$. Suppose $\mu \notin S(P, C, F)$. Since $\mu \in IR(P, C, F)$, there exists a pair (c, f) that blocks μ . Suppose $c \notin C'$, i.e., c is matched at some round k of MOD1. So at round k of MOD1, $\mu(c)$ is c 's first choice and thus $fP_c\mu(c)$ implies that f is already matched. That is, f is matched in the MOD1 phase as well, which implies that f was matched to his first choice at some round $k' < k$ of MOD1, and thus $\mu(f)P_fc$. So (c, f) cannot block μ , a contradiction. So, $c \in C'$. With a similar argument we also obtain $f \in F'$.

If $c \in C'$, then $\mu(c) \in F' \cup \{c\}$, and $f \in F'$ also implies $\mu(f) \in C' \cup \{f\}$. So (c, f) block $\mu|_{C' \cup F'}$ and thus $\mu|_{C' \cup F'} \notin S(P', C', F')$, a contradiction. So $\mu \in S(P, C, F)$.

We now show that $\mu = \gamma$. By Step 1, it suffices to show that for any $v \in C' \cup F'$, $\mu(v) = \gamma(v)$. Since γ is the candidate-optimal matching, for each $c \in C'$ and $\hat{\mu} \in S(P, C, F)$, $\gamma R_c \hat{\mu}$, and thus $\mu \in S(P, C, F)$ implies $\gamma R_c \mu$ for each $c \in C'$. Since $\gamma(C') \in C' \cup F'$, and $\mu(C') \in C' \cup F'$, $\gamma R_c \mu$ implies $\gamma R'_c \mu$. Likewise, for each $c \in C'$ and $\hat{\mu} \in S(P', C', F')$, $\mu R'_c \hat{\mu}$, and thus $\mu R'_c \gamma$ for each $c \in C'$. That is, $\gamma(c) = \mu(c)$ for each $c \in C'$. ■

⁴⁰Note that $v = c_2$ if c_2 's unique choice was matched in the first round.

References

- [1] Banerjee, S., H. Konishi and T. Sönmez. 1996. Core in a simple coalition formation game. *Social Choice and Welfare* **18**, 135–153.
- [2] Barrio, M. and M. Balinski. 2003. Admissions and Recruitment. *American Mathematical Monthly*, **110**, 386–399.
- [3] Balinski M. 2001. Les Enjeux du Recrutement. *Pour La Science*, **288** (octobre.)
- [4] Bouba-Olga O., Grossetti M. and Lavigne A. 2008. Le Localisme dans le Monde Académique: une Autre Approche. Réponse à Olivier Godechot and Alessandra Louvet. *La Vie Des Idées*.
- [5] Calsamiglia, C., G. Haeringer and F. Klijn. Forthcoming. Constrained School Choice: An Experimental Study. *American Economic Review*.
- [6] Dubins, L.E. and D.A. Freedman. 1981. Machiavelli and the Gale-Shapley Algorithm. *American Mathematical Monthly*, **88**, 485–494.
- [7] Gale, D. and L. Shapley. 1962. College Admission and the Stability of Marriage. *American Mathematical Monthly*, **69**, 9–15.
- [8] Gale, D. and M. Sotomayor. 1985. Ms Machiavelli and the stable matching problem. *American Mathematical Monthly*, **92**, 261–268.
- [9] Haeringer, G. and F. Klijn. 2009. Constrained School Choice. *Journal of Economic Theory*, **144**, 1921–1947.
- [10] Hitsch, G., Hortacsu, A. and Ariely, D (forthcoming). Matching and Sorting in Online Dating *American Economic Review*.
- [11] Kojima, F. and P.A Pathak. 2009. Incentives and Stability in Large Two-Sided Matching Markets. *American Economic Review*, **99**, 608–627.
- [12] McKinney, C.N., M. Niederle, and A.E. Roth. 2005. The Collapse of a Medical Labor Clearinghouse (and why such failures are rare). *American Economic Review*, **95**, 878–889.
- [13] Knuth, D. E., 1976. *Mariages Stables and Leurs Relations Avec d’Autres Problèmes Combinatoires*. Les Presses de l’Université de Montréal, Montreal, Québec., Introduction à l’analyse mathématique des algorithmes, Collection de la Chaire Aisenstadt.

- [14] Niederle, M. and L. Yariv. 2008. Matching through Decentralized Markets. Mimeo, Stanford University.
- [15] Niederle, M. and A.E. Roth. 2003. Unraveling Reduces Mobility in a Labor Market: Gastroenterology with and without a Centralized Match. *The Journal of Political Economy*, **111**, 1342–1352.
- [16] Roth, A. E. 1982. The Economics of Matching: Stability and Incentives. *Mathematics of Operation Research*. **7**, 617–628.
- [17] Roth, A. E. (1984) Misrepresentation and Stability in the Marriage Problem. *Journal of Economic Theory*, **34**, 383–387.
- [18] Roth, A.E. 1991. A Natural Experiment in the Organization of Entry Level Labor Markets: Regional Markets for New Physicians and Surgeons in the UK. *American Economic Review*, **81**, 415–440.
- [19] Roth, A .E. 2008. What have we learned from market design? *Economic Journal*, **118**, 285–310.
- [20] Roth, A. E. and E. Peranson. 1999. The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design. *American Economic Review*, **89**, 748–780.
- [21] Roth, A. E. and U.G. Rothblum. 1999. Truncation Strategies in Matching Markets — In Search of Advice for Participants. *Econometrica*, **67**, 21–43.
- [22] Roth, A.E. and M. Sotomayor. 1990. *Two-sided Matching: A study in Game-Theoretic Modeling and Analysis*, Econometric Society Monographs, No. 18. Cambridge University Press, Cambridge, England.
- [23] Roth, A. E. and X. Xing. 1997. Turnaround time and bottlenecks in market clearing: decentralized matching in the market for clinical psychologists *Journal of Political Economy*, **105**, 284–329.