

CAHIER DE RECHERCHE #1711E
Département de science économique
Faculté des sciences sociales
Université d'Ottawa

WORKING PAPER #1711E
Department of Economics
Faculty of Social Sciences
University of Ottawa

Labor Market Effects of Urban Riots: an Experimental Assessment

Emmanuel DUGUET^{*}, David GRAY[†],

Yannick L'HORTY[‡], Loïc du PARQUET[§] and Pascale PETIT^{**}

July 2017

^{*} Université Paris-Est Créteil, ERUDITE (EA 437), TEPP-CNRS (FR 3435), UPEC, UPEM, F-94010 Créteil cedex, France; e-mail: emmanuel.duguet@u-pec.fr.

[†] Department of Economics, University of Ottawa, Ottawa-Carleton School of Economics, 120 University Private, Ottawa, Ontario, Canada, K1N 6N5; e-mail: dmgray@uottawa.ca.

[‡] Université Paris-Est Marne la Vallée, ERUDITE (EA 437), TEPP-CNRS (FR 3435), UPEC, UPEM, F-77454 Marne-La-Vallée France; e-mail: yannick.lhorty@u-pem.fr.

[§] Université du Maine, GAINS et TEPP-CNRS, UFR Droit, Sciences Economiques, Gestion, Avenue Olivier Messiaen 72085 Le Mans cedex; e-mail: loic.duparquet@univ-lemans.fr.

^{**} Université Paris-Est Marne la Vallée, ERUDITE (EA 437), TEPP-CNRS (FR 3435), UPEC, UPEM, F-77454 Marne-La-Vallée France; e-mail: pascale.petit@u-pem.fr.

Abstract

We propose to measure the effects of urban riots on the labour market prospects of workers residing in affected areas through the channel of labour market discrimination based on locality. We investigate the case of the French riots of 2007, which were very geographically concentrated. The town of Villiers-le-Bel is selected as the treatment unit because it received a uniquely high degree of unfavourable exposure in the media. Two other towns serve control groups: i) Sarcelles, which is contiguous to Villiers-le-Bel, has a similar socio-economic-demographic profile, and did experience some rioting activity, and ii) Enghien-les-Bains, which is considered to be economically advantaged and did not experience rioting activity. Using the technique of correspondence testing, we are able to discern disparities in call-back rates for fictitious candidates who respond to actual job postings over four dimensions: gender, ethnic origin, locality of residence (advantaged vs. disadvantaged), and the degree of media exposure during the riots. We propose a new empirical approach to measure discrimination across several dimensions that integrates a set of relevant parameters into one unified system of equations with a compact and tidy structure. We decomposed the probability of receiving a callback for any candidate of given characteristics as a function of several parameters that are evaluated through difference-in-differences estimators. We find statistically significant negative effects of a pure media exposure effect. All other factors held constant, people residing in the area which received negative publicity were 3.2 percentage points less likely to receive a callback. The group of workers who tend to be the most associated with the riots, i.e. men of North African origin (at least in terms of perceptions), are the least affected by potential discrimination by region of residence, while women of French origin are the most affected.

Key words: Labour market discrimination, Call-back rates, Discrimination, Urban riots

JEL classification: C93, J71, R23

1. Introduction

In the aftermath of the wave of urban riots that occurred in the US in the late 1960s, a sizeable scientific literature dealing with the causes and consequences developed in the fields of political science and sociology. It was not until much later that economists such as Collins and Margo (2004,2007) applied the tools and the conceptual approach of economics to that phenomenon. Since that eventful period of widespread civil unrest, major riots erupted in Miami in 1980 and Los Angeles in 1992. Much more recently they broke out in the US cities of Ferguson, MO, Baltimore, MD, and Charlotte, NC in 2015 and 2016. France and the UK have also experienced major, sporadic civil unrest involving violence over the past three decades. The appertaining economic questions remain quite relevant, and in the case of the US disturbances, the economic impact has yet to play out.

A common theme of the literature on the economic impact of riots is that the subsequent effects affecting the demand side of product, labour, and real estate markets far outweigh the immediate negative supply-side impact in the form of physical reconstruction costs. Furthermore, the deleterious demand-side consequences are far more persistent. Some of these articles have generated empirical evidence of the existence of negative impacts on observed labour market outcomes in affected areas, such as higher unemployment rates, lower labour-force-participation rates, and lower employment rates. Nevertheless, less attention has been paid to the specific behavioural channels that could generate these unfavourable labor market impacts. One contributing cause is likely to be the emergence of depressed demand conditions in the labour markets in affected neighborhoods. Other contributing causes could work on the supply side through deleterious social network effects operating in these neighborhoods. Consequences such as ensuing out-migration, a deterioration in educational and social support services in the community, adverse peer effects, and other factors can hamper human capital development.

Our focus in this paper is on a specific demand-side behavioural mechanism that would weigh heavily on the labour market prospects of workers residing in affected areas, namely labour market discrimination based on locality. Such job applicants might be stigmatized by would-be employers who view that feature as a negative signal of productivity. Such a perception would have the effect of shifting the apparent wage-offer distribution facing the workers to the left and perhaps simultaneously lowering

the wage offer arrival rate. In addressing the topic of labour market discrimination based on the geographical area of residence, and to its degree of prior exposure to urban unrest, we borrow from the literatures on labour market discrimination as well as urban economics. To our knowledge, this paper is the first within the economic literature dealing with violent political conflicts and riots that explicitly involves labour market discrimination.

The methodology that we utilize is an experimental evaluation following the event of riots that occurred in France in 2007. More specifically, we adopt the “correspondence testing” technique that is similar to that employed by Bertrand and Mullainathan (2004) in a US context and Oreopoulos (2011) in the Canadian context. The major discriminating attribute in the application contained in the first paper is the race of job applicants, but their secondary focus serves as the primary attribute for our study, namely the worker’s neighborhood of residence. Our design allows us to achieve randomization as far as the unobservable influences are concerned and to condition on some critical observable variables, such as geographical distance and the attributes of gender and ethnic background. To our knowledge, our paper is the first within the literature on the economics of riots to employ an experimental methodology.

Another point of novelty is that our experiment involves two related but distinct treatments. The first one consists of the impact of the job applicant residing in an area where violent protests did occur but did not receive a wide coverage in the media. The second one consists of the impact of the job applicant residing in an area where violent protests also occurred but did not receive much coverage in the media. Whereas in the economic literature dealing with urban riots, impacts are typically identified by the severity or the scale of the civil unrest, we also include a measure of the degree of media coverage in our analysis. This is potentially important in that it could play a role in shaping the *perceptions* of would-be employers searching in an uncertain environment for information regarding the suitability and riskiness of unknown job candidates. The experiment is structured such that the two both of those different treatment effects are differentiated according to gender and to ethnicity, allowing us to condition our estimated effects on those two attributes as well as to estimate interacted effects.

We utilize an experimental data set that was first exploited in our earlier study (L'Horty *et al.*, 2011). We now extend that analysis along several dimensions. First, following Neumark (2012), who exploited the data set originally from Bertrand and Mullainathan (2004) in order to propose a new measure of discrimination, we construct a new empirical modeling of the probability of receiving a callback for any candidate. This probability is broken down into components associated with the candidate's characteristics and interactions between them. We are able to discern the disparities in callback rates originating from four sources: gender, ethnic origin, location, and the degree of media exposure following the 2007 riots. We estimate these parameters by applying the Asymptotic Least Squares (ALS) method to the probability components model. This technique can also be interpreted as a generalization of the difference-in-differences estimators to the over-identified case. Our model is tidy, compact, and unified into one system which can be estimated in one command without generating excessive empirical detail. Second, we couch it within the conceptual framework of the economics of riots and political conflict. We conclude that riots can have significant consequences on subsequent labor market outcomes, and that these consequences differ between demographic groups.

2. Empirical strategy: 'a tale of two cities'

The riots were both intense and very geographically concentrated. They were triggered immediately after the accidental deaths of two adolescent boys who were both sons of African immigrants. They were killed when their motorcycles hit a police car in November of 2007 in the town of Villiers-le-Bel, which is situated in the far Northern suburbs of Paris. Over the next two nights, hundreds of masked youths participated in angry protests incinerating cars, stores, and public buildings, including notably a school and a library. There were violent attacks against police officers, a few of which (for the first time in modern French history) involved firearms.¹

¹ These were hardly the first wave of riots which occurred in France. Beginning in the late 1970s and continuing throughout the 1980s, social movements involving violence occurred in suburban neighborhoods of Lyon and elsewhere. These conflicts spurred the development and implementation of a new set of urban policy measures, including the creation of the designated "zones d'éducation prioritaires (ZEP)" (priority zones for education). The wave of riots that broke out in October of 1990 led to the creation of a totally new Ministry of Cities. Despite this attention from governments, similar incidents recurred throughout the 1990s. Those events gave rise to new policy responses by the government, with the creation of the "zones urbaines sensibles (ZUS)" (sensitive urban zones) in 1996. In the new millennium, widespread riots broke out in many French cities in October of 2005. When the

In order to investigate the impact of the riots, we selected two towns that are located in the department of Val d’Oise (analogous to a suburban county in the USA), to serve as control groups for Villiers-le-Bel: Sarcelles and Enghien-les-Bains. These three towns are equidistant from Paris in terms of transportation, which allows us to abstract from any effect tied to the physical distance between the places of residence and employment (see map 1). Furthermore, Sarcelles is contiguous to our treatment town of Villiers-le-Bel, and those two towns have very similar socio-economic-demographic profiles. Both are considered to be disadvantaged, whereas the other control-group town, Enghien-les-Bains, is viewed as advantaged. The unemployment rate and the percentage of adults without any diplomas whatsoever are high in both Sarcelles and Villiers-le-Bel but low in Enghien-les-Bains. More than 60 % of the inhabitants of Sarcelles and Villiers-le-Bel reside in designated neighborhoods, in contrast to 0 % in Enghien-les-Bains. In fact, the level of per-capita income in Enghien-les-Bains is one of the highest in the department of Val d’Oise.

[map 1 here]

The identifying hypothesis for our empirical analysis is that the primary difference between Sarcelles and Villiers-le-Bel involves the degree of media coverage of the riots. As mentioned above, Enghien-les-Bains is an advantaged area, whereas the towns of Villiers-le-Bel and Sarcelles are similar in terms of socio-economic-demographic composition, the living standards of their inhabitants, and the distance to employment. We attribute the estimated differences in the testing outcomes between these two towns to effects identified by variation in the level of media coverage. That choice for a treatment indicator is based on evidence that is presented in graph 1, which displays times series data on the number of “Google hits” that were obtained since the year 2005 when the names of each town were inputted. All three curves reflect indices which are benchmarked such that the value for Villiers-le-Bel is equal to 100 in the month of November 2007, which corresponds to the peak in the rioting activity. Note that this remarkable spike is mostly limited to November, although there are minor upticks appearing in the following two months reflecting subsequent arrests. Shortly thereafter, the series for “Google hits” regarding Villiers-le-Bel falls back to the level that existed before the riots. Neither

riots of 2007 first broke out, many feared that they would spread throughout France, but their scope remained limited to certain municipalities located in the Northern suburbs of Paris.

Sarcelles nor Enghien-les Bains exhibit any such intertemporal pattern. The pattern for the latter town is very stable over the entire interval. Perhaps surprisingly, the level for the relatively well-off town of Enghien-les-Bains is very close in magnitude to that of the economically disadvantaged town of Villiers-le-Bel, with the exception of the short-lived spike.

We also indicate on the timeline of Figure 1 the period during which we conducted the experiment and collected our data, which extended from mid-December 2008 to the end of January 2009. This interval is set a little over a year after the riots occurred, by which time the degree of media interest regarding these three towns had reached normal (i.e. low) levels.

[figure 1 here]

3. Survey of the Literature

3.1 The interface between labour market discrimination and the place of residence

There exists a substantial and solid body of quantitative analysis of labour market statistics in France, at a detailed level of aggregation, which indicates the existence of sharp discrepancies in outcomes such as unemployment rates and jobless/unemployment durations, such as papers by Gobillon *et al.* (2011), Gaschet and Gaussier (2004), and Duguet *et al.* (2008). These empirical patterns led some authors to consider the possibility that labour market discrimination could be a contributing factor. Duguet *et al.* (2009) was one of the first papers to examine the hypothesis that employers might utilize the place of residence as a criterion for hiring in addition to the ‘usual suspects’ of ethnic origin, gender, and race. Indeed, the perception that discrimination based on the place of residence was sufficiently widespread that the government passed a law in 2014 (No. 2014-173) that explicitly outlawed the practice.

First we discuss the spatial dimension of employment outcomes. According to Gobillon *et al.* (2007), there are a number of channels through which the location of the residence of the worker can influence the probability of gaining employment, any combination of which could generate a spatial pattern. Perhaps the most straightforward of them involves ‘local skill mismatch’, whereby within the

confines of a spatial unit, either there is no demand on the part of employers for local workers possessing a given skill set, or there is no supply of workers possessing the skill sets sought by employers.

An alternative approach of ‘spatial mismatch’ emphasizes the role of physical distance between the location of attractive jobs and the location of the job searchers, as well as the adequacy of the means of transportation linking the two spatial units. From a strictly cost-benefit perspective, the greater the distance, the higher the transportation costs, and the less attractive any given job offer becomes. Reinforcing the explicit costs of transportation are implicit costs associated with the existence of cheaper housing typically being available in distant residential neighborhoods, which further lower the returns to job search. Applications of this approach include papers by Gauthier and Zenou (2010) and Patacchini and Zenou (2005).

Another behavioural channel through which access to employment can exhibit a strong spatial pattern is through neighborhood and social network effects which transmit information regarding job offers. The basic thrust of this approach is that in disadvantaged neighborhoods, these networks tend not to be extensive and not to convey useful information regarding job market opportunities. The nefarious social network effects can be augmented by peer effects that might hinder human capital formation. According to either of these two mechanisms, the *collective* attributes of individuals within a spatial unit are thought to hamper job search activities and to lower the probability of receiving an offer for an attractive job.²

The conceptual underpinning for the observed empirical, spatial pattern of labour market outcomes on which we focus is labour market discrimination associated with preferences of employers with respect to factors that are unrelated to productive attributes. There are at least three conceptual approaches that are contained in the economic literature regarding labour market discrimination. According to the earliest approach set out by Becker (1957,1971), its sources lie in the exogenous aversions of employers, co-workers, or the clientele - typically with respect to a visible attribute.

² This approach is surveyed by Ioannides and Datcher-Lourey (2004) and Galster (2012). Two papers by Hellerstein *et al.* (2008,2014) consist of applications to urban labor markets in the USA with a heavy concentration of African-American workers. Fontaine (2006) applies this approach to the French labor market.

According to the more recent alternative approach of statistical discrimination attributed to Phelps (1972) and Arrow (1973), given imperfect information regarding candidates' productivity, firms resort to signaling in their recruitment procedures. In our particular context, the place of residence serves as a signal for desirable (or sometimes undesirable and risky) yet unobservable attributes based on employers' perceptions regarding the town where the potential workers reside. The probability of an applicant being initially assigned a high score for expected performance is a direct function of the perceived average score for inhabitants of that spatial unit.

"Red-lining" is a third conceptual approach of labour market discrimination that not only can manifest itself empirically in a spatial pattern, but is explicitly oriented around spatial contours, as set out in Zenou (2002, 2006). As is the case with statistical discrimination, employers cannot accurately observe the productivity of job candidates, and seek signals to fill that informational void. The hiring criterion becomes the physical distance between the workers' residences and the workplace, or perhaps some composite measure of the costs of transportation. The employer draws an imaginary 'red-line' beyond which he/she perceives that the transportation costs for a potential recruit are sufficiently high to deter any would-be productive employee from applying, while unproductive employees beyond that boundary might be less deterred. As a consequence, all workers residing on the "wrong side of the tracks" are less likely to receive offers.

3.2 The economic impact of riots and political conflict

DiPasquale and Glaeser (1998) is an early and influential study with a descriptive title: "The Los Angeles Riot and the economics of urban unrest". It is more focused on the causes of riots rather than on their impact. Their data set contains information on international riots, cross-city US evidence of race riots from the 1990s, as well as the Los Angeles riots in 1992. One major finding is that holding city size constant, riot size as well as riot intensity are most closely linked to the size of the non-white community. Overall they assert that the neo-classical view of rioting - in which individuals' benefits and costs play an important role - is more applicable than approaches stressing the role of community and social conditions.

The seminal, influential piece by Collins and Margo (2004) lays out the basic approach for discerning the impact of riots on economic outcomes. It is based on a difference-in-difference methodology that compares blacks' labour market outcomes in cities that had severe riots with their counterparts in cities which did not experience severe riots during the period of the late 1960s. Riots are viewed as a negative shock to labour market conditions. The essence of the medium and long-run impacts consists of depressed demand conditions in both input and output markets. The estimated riot effect is identified by variation in the degree of severity of the disturbances across cities derived from census data from 1960, 1970, and 1980.

The first part of their empirical analysis assesses the impacts on the labour market outcomes of income, employment, and unemployment measured at the city level. The authors address endogeneity concerns by arguing heuristically that variation in riot severity across locations was highly unpredictable and spontaneous, and hence orthogonal to the *ex ante* socio-economic conditions prevailing in a community. They are most interested in endogenously-propagated indirect effects that unfold over a longer period taking the form of a vicious circle of "knock-on" effects. The empirical findings indicate that the riots had negative effects on blacks' incomes and employment rates that were economically significant, especially in the longer term. The ordinary least squares estimates yields underestimates, indicating that the riot intensity regressor is indeed endogenous.

The second component of Collins and Margo (2004) involves individual-level data on outcome variables of income and employment status. They find evidence of a declining employment rate generated by both a fall in the labour force participation rate as well as by a rise in the unemployment rate. Their empirical analysis abstracts from spillover or out-migration effects, as those who migrate out of the affected cities are not reflected in the data.

A later piece by Collins and Margo (2007) focuses on the impact of the 1960s riots on property values. This article measures the riots' impact on the value of central city residential property, and especially with respect to black-owned property. Although they do not quantify perceptions (in this context they refer to 'neighborhood amenities'), they mention that they are an important demand-side influence. It is not the initial physical damage *per se* that negatively affects property values, but rather

people's perceptions of the event's severity and implications. Their results indicate that there were relative losses of population and property value in tracts that were directly affected by riots compared with other tracts in the same cities. Their interpretation for the endogenous and persistent decline is that negative external effects in the context of racially segregated housing markets were at work. Proximity to riots is viewed as a harbinger of persistently higher risk, which gives rise to a sustained negative shift in perceived local amenities.³

Abadie and Dermisi (2008) deal with a related topic, namely the economic impact of terrorism rather than of urban riots. The approach and methodology are similar. The outcome variable for their study is real-estate vacancy rates in the central business district of Chicago before and after the terrorist attacks of September 2001. Their empirical results indicate that office vacancy rates increased in the wake of the attacks. Buildings with characteristics that caused them to be perceived as vulnerable to attacks experienced a particularly severe deterioration in real estate market outcomes.

Braakman (2012) is a more recent study regarding the impact of the London riots of 2011. While his primary focus is on subsequent criminal behavior, he does investigate the impact of riots on the local unemployment rate as a proxy for changes in economic conditions in affected areas. He concludes that the riots increased unemployment statistics; specifically, the levels of unemployment insurance claimants of both short and longer-term duration increased by sizeable magnitudes. Regarding the unemployment estimates, the author states that "...it seems possible that the riots led to some stigmatization of individuals from those areas in the labour market, which should also lower labour market prospects for people from affected areas" (p. 14). This interpretation is quite germane to our paper, as both papers posit that negative labour market outcomes stem from labour market discrimination with a spatial dimension.⁴

³ Matheson and Baade's (2004) brief piece deals with the economic impact of only the Los Angeles riot of 1992. A unique feature is that they employ a strictly time-series analysis to the variable of taxable sales. They determine that the riots had a long-lasting, negative impact.

⁴ In Braakman's work it the number of workers claiming the basic UI benefit called the "Jobseeker's allowance (JSA)", which is reported over a pretty narrow census tract. In our study it is the response (or lack thereof) of an application to a posted job offer.

The well-known paper by Besley and Mueller (2012) focuses not on the impact of riots *per se*, but rather on the impact of continual political violence and subsequently its cessation - the so-called “peace dividend” that occurred in Northern Ireland. Following the approach of Collins and Margo (2007), they exploit variation in the degree of violence within Northern Ireland to assess empirically the impact on house prices. This particular piece is more grounded in economic theory than is the case for most articles in this literature, as they develop a formal model linking house prices to violence. Their data set is also particularly rich, as it has a longitudinal structure containing spatial variation as well as temporal variation. The results indicate that there tend to be larger peace dividends in regions where violence was more severe and more persistent. Another interesting finding is evidence on spillover effects across areas, suggesting that some kind of general equilibrium response to violence involving migration might be at work.⁵

We cite Huet-Vaughn (2015) because it shares the attribute of using French data, albeit from an earlier period from 1980-1995. The context of his work is strikes, political demonstrations, and low-grade violence. Unlike our investigation, he does not focus on their economic impact, but rather on its political efficacy from a public choice perspective. One point of novelty is that he exploits a data set containing data with a daily frequency. The outcome variable is the likelihood that violent protest causes policy concessions. He discerns a significant and negative relationship between property destruction associated with the protests and the chance of near-term success in changing policy.

What our paper shares with most of the papers mentioned above is the objective of assessing the economic impact of protest or riot violence in the causal sense. Another common point is that perceptions of either employers or property buyers can play a role in contributing to that impact. Following those studies, we adopt a spatial perspective in order to achieve empirical identification.

⁵ A study conducted by Arbel *et al.* (2010) based on Israeli data is in the same spirit as Besley and Mueller (2012). These authors analyse changes in house prices in the Gilo neighborhood of Jerusalem following the Second Palestinian Intifada in the year 2000, which was subjected to sporadic gunfire from a nearby Palestinian village. They find that the shooting events caused a short-run 12% reduction in values, but these values tended to recover within 18 months. Their primary conclusion is that there exists substantial negative residential property value effects stemming from terrorism, and that households living in conflict areas would benefit from conflict resolution.

Those papers deal with an assortment of economic outcomes, including home values, commercial building vacancy rates, and labour market statistics such as unemployment, employment, and participation rates – usually at a fairly aggregated level.

A point of difference is that we cannot model the same outcomes as Collins and Margo and Braakman do – i.e. labour market outcomes for visible minorities – because French data sets are prohibited by law from having any information regarding race or ethnicity on actual individuals. We focus our attention on a more specific and individualized outcome of access to employment opportunities, namely potential labour market discrimination on the part of employers. While we cannot use real candidates as subjects displaying a certain ethnicity, we do utilize fictitious candidates whose files contain markers for ethnicity. Our methodological approach differs from that which has typically been employed in this literature, as we carry out an event-based, experimental investigation, while the studies cited above all utilize non-experimental econometric methodologies such as fixed-effects or difference-in-difference strategies in efforts to identify the economic effects of riots.

3.3 The experimental method

Our research design aims to control for any effect related to the physical distance between the places of residence and employment in order to focus directly on locational effects. By the selection of towns, this distance is roughly equivalent for all subjects, thus permitting us to exclude the pure spatial mismatch approach and red-lining (based solely on distance) as being relevant. We judiciously choose residential areas that are contiguous, rendering any line of demarcation between job applicants residing in the two regions infeasible for the employer in search of information on unobserved attributes. While we are able to disentangle the effects of physical distance from locational effects, we are unable to discern the sources of employers' preferences or aversions regarding workers' locations, i.e. whether they are shaped by classical discrimination or by information-based discrimination. Although we lean towards the latter interpretation, we cannot rule out the proposition of pure aversion based on our empirical analysis.

The field-experimental methodology that we employ is also called “correspondence testing”. It utilizes fictitious applications whose human capital attributes are presented as identical across groups. Every effort is made to eliminate differences in average characteristics across groups, barring the criterion on which the potential discrimination is based. As pointed out in Neumark (2012), identification is based on the assumption that employers believe that there are no average differences in observable characteristics across groups. The only outcome variable is the event of a callback, thus obviating the possibility that individual job applicants can influence the outcome of the experiment at stages subsequent to the initial point of contact. While it has been applied in the context of discrimination by researchers in other social science disciplines for decades, economists have not adopted this approach until relatively recently. An influential paper by Riach and Rich (2002) encouraged economic researchers to utilize it in their studies regarding discrimination.⁶

Bertrand and Mullainathan (2004) published their influential application of labour market discrimination based on race and ethnicity set in Chicago and Boston. We structure our research design and protocols along their lines. We also follow their approach in modeling the impact of neighborhood of residence in addition to the impact of ethnic origin, as well as the interacted effect. In Bertrand and Mullainathan, however, the locality effect is interpreted as an indicator of the socio-economic status of the applicant, whereas in our study we interpret it as an indicator of the degree of violence that ensued during the riots. They find that applicants with ‘white-sounding’ names are 50 percent more likely to receive a callback than is the case for candidates with ‘African-American sounding’ names. For our purposes an interesting result of their study is that while the trait of residing in a wealthier neighborhood increases callback rates, that advantage does not differ between the two racial groups.

Tunstall *et al.* (2014) is a British study utilizing the correspondence testing technique involving fictitious candidates for actual job postings (requiring relatively low skill levels) in England and Wales in search of discriminatory treatment against applicants residing in neighbourhoods with ‘poor and bland reputations’, which they label ‘postal code discrimination’. The outcome variable for the experimental

⁶ Surveying results pertaining to ten different countries involving several types of markets, those authors conclude that there does exist pervasive discrimination against women and non-whites.

component of their article is the event of a call-back at the initial stage. Unlike some other articles in this literature, they do not reject the null hypothesis of equal treatments of the two classes of candidates.

Turning to the French literature, among the first papers employing this methodology in order to address the question of job discrimination against workers of immigrant origin is by Duguet *et al.* (2010) in an application involving accountants. While those authors did find evidence of discrimination against candidates according to ethnic origin, they discerned no difference in treatment according to neighborhood of residence. The study authored by L'Horty *et al.* (2011) was one of the first French studies to apply this methodology with an eye towards modeling the effects of locality as well as ethnic origin. They find that only women are affected. Irrespective of their ethnic origin, they are called back at a lower rate if they reside in a disadvantaged neighborhood. In the case of women with French-sounding names, it appears that those residing in a neighborhood that was subjected to a lot of unfavourable media coverage are disadvantaged. Petit *et al.* (2014) involves a field experiment conducted in the greater Paris area utilizing correspondence testing that varied the spatial contours of the workers' localities. The research design is structured around six neighborhoods in total comprising three socio-economic levels within either a large "department" (like an urbanized county) or an "arrondissement" of Paris (like a large city ward). The neighborhoods are selected to reflect a continuum of socio-economic levels. They find that "a good address can triple the chances of being invited to a job interview" (p. 545), and that candidates residing in a disadvantaged county face a greater handicap than those residing in a disadvantaged ward.

4. Data collection

The basic empirical strategy for any field experiment which seeks to discern discriminatory behavior on the part of would-be employers is to ensure that all candidates have identical productivity-related attributes, as indicated in explicit terms on their applications. These specifications should capture observable human capital attributes such as formal training and education, on-the-job-training, and cognitive skills. Factors such as non-cognitive skills cannot be observed by would-be employers at this initial stage of the hiring process, and are thus irrelevant. Fluctuations in the exogenous variables that are the focus of this study, namely gender, ethnic background, and the town of residence are generated

by design. While other aspects of their CVs, particularly with respect to how they are viewed by the recruiting firms, necessarily vary across candidates, ideally the expected values of these variables are equal to zero via the process of randomization. This means that the conditional independence condition should apply if the selection of job offers in our sample is representative of the population of firms which hire workers in this occupation in the greater Paris region.

In our earlier study (L'Horty *et al.* 2011), we conduct robustness checks in order to assess the validity of that critical assumption. There exists a slew of other control variables for which we have data that are unrelated to the job applicant but potentially have some influence on our outcome variable of receiving a call back. They are grouped into five categories: i) variables related to the testing procedure, ii) variables related to the job offer, iii) characteristics of the firm, iv) characteristics of the towns where the position is located, and v) characteristics related to the potential commute for that candidate and the job in question. When these (approximately) 25 supplementary regressors are included in the basic specifications from which the estimated coefficients of discrimination are derived, the results are qualitatively identical and quantitatively very similar. We interpret these findings as evidence that these variables are not highly correlated with the regressors included in the trunk of the estimating equations. The omission of these variables from the specification contained in this paper, therefore, should not impart biases for the estimates of the various components of discrepancies in call-back rates.

We now turn to a detailed description of the protocol. The sampling took place over the six-week period from mid-December 2008 to the end of January 2009. This window occurred during the Great Recession, when employment was falling at a rapid pace. A total of 307 positions were selected for the sending of our responses for the fictitious candidates. Because the experiment involves two ethnic origins (French and Moroccan), three towns (Enghien les Bains, Sarcelles, and Villiers-le-bel), and two genders, a total of 12 distinct applications were sent out to each offered post. The total sample size is thus 3,684 ($12 * 307$).

As far as the applications are concerned, the ethnic origins are not explicitly stated on the CVs, nor are any photographs included. Such practices would not only violate French labour law, but they would also invalidate the experiment. Instead, as is common in the literature, the origins are indicated

by common surnames and given names that are either French-sounding (e.g. Emilie Durand) or Moroccan-sounding (e.g. Yasmina Brahim). All of them are 25 year-old singles of French nationality. The CVs indicate that they have a driver's license, which is often viewed as an indicator of mobility. While by design two-thirds of our profiles reside in towns having relatively low socio-economic status (Sarcelles and Villiers-le-bel), none of them declare their residence in a neighborhood classified as disadvantaged and targeted for governmental assistance.⁷

In regards to their productive characteristics, we solicited the advice of professional recruiters within the occupation in order to validate that the credentials which we specify are realistic, credible, and relevant. They also ascertained that the CVs were devoid of peculiarities, salient features, and potential "red flags". All candidates have the same academic background commencing with a high-school diploma with a scientific orientation ("BAC – science"). That was followed by the successful completion of a three-year undergraduate degree in information technology ("Licence – informatique"), which in turn was followed by the successful completion of a two-year master's degree in IT ("M1 plus M2 en informatique"). The nomenclature for the entire academic itinerary is Bac + 5 ("high school plus five").⁸ They all did internships in conjunction with both years of their master's degrees, which are similar to co-op programs in North America. After completing the second placement, they were hired on as full-time, salaried workers at the same firm where they worked during their second internship. Thereafter they have two years of work experience in IT design and development at the same firm. They all list the same set of more specific skills: programming, operating systems, web development, data bases, project management, and protocols. They are still employed at this firm at the time of the application, and there are no gaps in the timeline that could be attributed to a bout of unemployment or absence from the labour force.

The occupation of information technology was chosen for two reasons. First, we sought a labour market that was sufficiently tight such that a discriminating firm would incur high costs. Furthermore,

⁷ This designations, made for public policy purposes, were labelled "Zone urbaine sensible", or ZUS. It is likely that most employers were aware of the broad outlines of this governmental program, and some might have been aware of the boundaries of these neighborhoods.

⁸ These master's level courses are in different universities in the Ile-de-France (Greater Paris) Region, but they are of similar reputation and confer computer skills comparable to the twelve candidates.

given such labour market conditions, there would not be a wide margin for refusals based on other reasons that would confound discriminatory hiring. The indicator that we adopted to ensure this condition was a high exit rate from unemployment for workers in the occupation. Second, we sought a labour market that was thick and of large scale such that the experiment would not be detected by the recruiters. We wanted our batch of applications to be diluted in the labour market in order to reduce the chances of an occurrence of the Hawthorne effect. To this end, we selected an occupation characterized by a large stock of unemployed workers giving rise to a large recruitment pool. Data for these two variables are drawn from Pôle Emploi administrative data over the period from 2003-2006.⁹

Turning to the selection of the job offers, we responded to postings for full-time jobs from three different categories of employers: ‘head hunters’, downstream users of IT, and IT firms. All of our applications were filed on the same day that the offer was posted. The matching from our set of CVs to the job posting was done by algorithm in order to ensure randomization with respect to specific features, such as the particular university attended or the particular firm that currently employed the candidate. The sequence in which the twelve CVs were sent out alternated across job postings.

We aim to have a representative sample of firms that recruit from this labour market without an undue concentration of anomalous firms. To this end we calculated a histogram of the number of *positive* responses registered per offer over the nodes of 1 through 12, with 12 being the maximum possible number of positive responses. If our sample were composed of a significant number of atypical firms, we would not expect to receive multiple positive responses at a high frequency. We did not discern that empirical pattern, however. Among the job offers for which we received at least one positive response, there were two or more responses in 83 % of the cases.

As mentioned above, we are concerned about the risk of detection, given that we send out 12 fake CVs to every job offer, which is more than what is typical in correspondence experiments. According to the administrative database for the “pole d’emploi”, in the fourth quarter of 2008, for every

⁹ The indicator for the first criterion is the Pole Employment Index for Tension, which divides the number of Registered Job Offers (*Offres d’Emploi Enregistrées* –OEE) by the number of End-of-Month Job Applications (*Demandes d’Emploi en Fin de Mois* - DEFM) at the level of each detailed occupation. The indicator for the second criterion is the sum of OEE and DEFM.

IT job posted in France, the average number of responses across all 95 departments was 32. The figure for the Paris metropolitan region is far greater, and so hopefully our particular group of 12 CVs would not appear to be conspicuous. To address this issue, we calculated a histogram of the success rate (or the call-back rate) as a function of the sequence in which the 12 applications were sent out. If we were sending out too many applications, then we would expect to see a decreasing pattern indicative of diminishing returns – the ones that were sent last receiving fewer call-backs. The empirical pattern that we do discern, however, appears to be slightly uniform.¹⁰

Table 1 shows the success rates for each profile for the same vacancies. Of the total number of applications sent, more than half (52.1%) received at least one favorable response, which reflects the high level of tension existing in this labor market (BAC + 5 computer scientists). Among French candidates, men residing in Sarcelles and women residing in Enghien and Sarcelles had the highest success rates (19.9%, 22.5% and 22.1%, respectively). Conversely, men who reside in Enghien or Villiers-le-Bel have lower success rates (16.9% and 18.6% respectively), as well as women in Villiers-le-Bel (17.9%). Among the candidates of Moroccan origin, women residing in Enghien had the highest success rate (19.5%), as did men in Enghien or Sarcelles (18.6% and 19.2%, respectively). Candidates of Moroccan origin with the lowest call-back rates are men residing in Villiers-le-Bel (17.3%) and women in Sarcelles (13.7%) and in Villiers-le-Bel (15%).

[Here Table 1]

5. The Empirical model

We seek to discern disparities in call-back rates over four dimensions: gender, ethnic origin, location, and degree of media exposure. To this end we construct a probability components model for which the unit of observation is the cell into which candidates are grouped according to those criteria. The estimates for the disparities are generated as simple linear functions of the estimates of components and the interactions between them. Two of the components, origin and gender, reflect directly on the job seekers, while the other two pertain directly to the location. The first of these captures whether or

¹⁰ Details for both of these diagnostic histograms are provided in l'Horty *et al.* (2011).

not the town is advantaged or disadvantaged, while the second one reflects the media exposure or reputation effect. In order to express the probability that a worker characterized by ethnic origin \mathbf{o} , gender \mathbf{g} , location \mathbf{l} , and media exposure \mathbf{m} receives a positive response from job posting \mathbf{i} , we specify the following general equation with this nomenclature:

$$P(\mathbf{o},\mathbf{g},\mathbf{l},\mathbf{m}) = O(\mathbf{o}) + G(\mathbf{g}) + L(\mathbf{l}) + M(\mathbf{m}) + OG(\mathbf{o},\mathbf{g}) + OL(\mathbf{o},\mathbf{l}) + OM(\mathbf{o},\mathbf{m}) + GL(\mathbf{g},\mathbf{l}) + GM(\mathbf{g},\mathbf{m}) + u_i$$

The first four terms denote functions representing the first-order or direct effects: O for origin, G for gender, L for location, and M for media exposure. The remaining five terms denote functions representing the joint or the interacted effects: origin crossed with gender (OG), origin crossed with location (OL), origin crossed with media exposure (OM), gender crossed with location (GL), and gender crossed with media exposure (GM). There are two codes for ethnic origin: French (f) and Moroccan (m). The codes for the gender attribute are (m) for men and (w) for women. The codes for the location attribute are advantaged (a) and disadvantaged (d), and for the degree of media exposure, they are non-exposed (n) and exposed (e).

The disturbance term u_i captures the job offer unobserved heterogeneity, which is specific to each job posting and is likely to be correlated with all the other components of the model. In order to estimate the discrimination components, we will first eliminate the unobserved heterogeneity across jobs. This is achieved through differencing across the 12 observations (for each job offer) that correspond to each candidate relative to a benchmark candidate, as explained below.

The three cities where the experiment was conducted have the following characteristics:

- Enghien-les-Bains: advantaged (a) and no media exposure (n). The total locational effect for this town is equal to: $L(a) + M(n)$
- Sarcelles: disadvantaged (d) and no media exposure (n). The total locational effect for this town is equal to: $L(d) + M(n)$
- Villiers-le-Bel: disadvantaged (d) and considerable media exposure (e) during the 2007 riots. The total locational effect for this town is $L(d) + M(e)$

The choice of the joint effects on which we focus is based on the following assertions. First, women of foreign origin could experience a degree of discrimination that is greater than the sum of the gender and origin discrimination coefficients (i.e. the OG effect, as treated in Petit *et al.* (2013)).

Second, the media focused their coverage of the riots on events involving men of foreign origin residing in disadvantaged localities. Given this situation, the degree of discrimination that they face could be compounded by including the interacted effects of OL, OM, GL, and GM.

Given that the above equation is a fully saturated model, it is clear that it contains too many parameters to allow for an estimation of all of them. It is necessary to set identification constraints - some of which can be tested. Some of these constraints are mere normalizations or conventions and thus do not raise concerns regarding generality. For instance, since we are interested in estimating the degree of gender discrimination, the quantity $G(m) - G(w)$ will suffice. The individual components involving G need not be estimated separately; only one of these parameters needs to be estimated. On the other hand, there are instances in which constraints must be imposed on the data in order to achieve identification, and it is not possible to test for the validity of all of them. Globally, however, since we have more constraints than are required for the purposes of estimation, our model is over-identified, and we can perform suitable over-identification tests.

The data set is structured around twelve distinct candidate profiles, which we sent out to actual posted job offers. The summary statistics regarding the call-back proportions are listed in Table 1, and they serve as our base values for estimating the following derived quantities of interest:

- gender discrimination : $\Delta G = G(m) - G(w)$
- origin discrimination : $\Delta O = O(f) - O(m)$
- location discrimination : $\Delta L = L(a) - L(d)$
- media exposure discrimination : $\Delta M = M(e) - M(n)$
- five interacted effects that reinforce or diminish the pure discrimination coefficients
 - women of Moroccan origin: $GO(w,m)$
 - Moroccan origin crossed with disadvantaged location: $OL(m,d)$
 - Moroccan origin crossed with media exposure: $OM(m,e)$

- men residing in a disadvantaged location: $GL(m,d)$
- men residing in an area exposed to the media: $GM(m,e)$.

Table 2 displays the structure of the probability components for each of the twelve distinct candidate profiles that are the subjects of our study. The total probability of receiving a callback for any candidate of given characteristics is equal to the sum all of the components which appear in that same row. For each one, the order is (origin, gender, neighborhood type, degree of media exposure), and the nomenclature is as follows: **f** for French, **m** for Moroccan; **w** for woman, **m** for male; **a** for advantaged, **d** for disadvantaged; **n** for no media exposure, and **e** for exposure. To give an example from the first row, the probability of a woman of French origin residing in Enghien obtaining an interview is equal to $P(f,w,a,n) = O(f) + G(w) + L(a) + M(n)$.

[Here Table 2]

Based on the structure presented in Table 2, we are able to evaluate the parameters of interest employing either simple difference estimators or difference-in-differences estimators. The data values consist of the realized percentage of a group of candidates that received a call back. In order to estimate discrimination coefficients, we take the difference with our selected baseline case: French-origin males living in Enghien (i.e. advantaged French men with no media exposure - the case of (f,m,a,n)). The relevant quantities consist of the differences in the success rates between the different groups and that baseline group. There are a total of eleven independent probability differences that can be related to the nine structural parameters as enumerated below.

- 1) $\Delta(f,w,a,n) = p(f,m,a,n) - p(f,w,a,n) = G(m) - G(w) = \Delta G$ (French woman, advantaged town, no media exposure gives the pure gender effect).
- 2) $\Delta(m,w,a,n) = p(f,m,a,n) - p(m,w,a,n) = O(f) - O(m) + G(m) - G(w) - OG(m,w) = \Delta O + \Delta G - OG(m,w)$ (Moroccan woman, advantaged town, no media exposure gives the gender effect compounded with the ethnicity effect).
- 3) $\Delta(m,m,a,n) = p(f,m,a,n) - p(m,m,a,n) = O(f) - O(m) = \Delta O$ (Moroccan man, advantaged town, no media exposure gives the pure ethnicity effect)
- 4) $\Delta(f,w,d,n) = p(f,m,a,n) - p(f,w,d,n) = G(m) - G(w) + L(a) - L(d) = \Delta G + \Delta L$ (French women, disadvantaged location gives gender effect compounded with location effect)

- 5) $\Delta(f,m,d,n) = p(f,m,a,n) - p(f,m,d,n) = L(a) - L(d) - GL(m,d) = \Delta L - GL(m,d)$ (French man, disadvantaged location gives ethnicity effect compounded with the location effect)
- 6) $\Delta(m,w,d,n) = p(f,m,a,n) - p(m,w,d,n) = O(f) - O(m) + G(m) - G(w) + L(a) - L(d) - OG(m,w) - OL(m,d) = \Delta O + \Delta G + \Delta L - OG(m,w) - OL(m,d)$ (Moroccan woman, disadvantaged location gives ethnicity effect compounded with the location effect and the gender effect)
- 7) $\Delta(m,m,d,n) = p(f,m,a,n) - p(m,m,d,n) = O(f) - O(m) + L(a) - L(d) - OL(m,d) - GL(m,d) = \Delta O + \Delta L - OL(m,d) - GL(m,d)$ Moroccan man, disadvantaged location gives ethnicity effect interacted with the location effect
- 8) $\Delta(f,w,d,e) = p(f,m,a,n) - p(f,w,d,e) = G(m) - G(w) + L(a) - L(d) + M(n) - M(e) = \Delta G + \Delta L + \Delta M$ (French women, disadvantaged location, media exposure gives gender effect compounded with the location effect compounded with the media exposure effect)
- 9) $\Delta(f,m,d,e) = p(f,m,a,n) - p(f,m,d,e) = L(a) - L(d) + M(n) - M(e) - GL(m,d) - GM(m,e) = \Delta L + \Delta M - GL(m,d) - GM(m,e)$ (French man, disadvantaged location, media exposure gives location effect compounded with the media exposure effect)
- 10) $\Delta(m,w,d,e) = p(f,m,a,n) - p(m,w,d,e) = O(f) - O(m) + G(m) - G(w) + L(a) - L(d) + M(n) - M(e) - OG(m,w) - OL(m,d) - OM(m,e) = \Delta O + \Delta G + \Delta L + \Delta M - OG(m,w) - OL(m,d) - OM(m,e)$ (Moroccan woman, disadvantaged location, media exposure gives ethnic effect compounded with gender effect compounded with location effect compounded with the media exposure effect)
- 11) $\Delta(m,m,d,e) = p(f,m,a,n) - p(m,m,d,e) = O(f) - O(m) + L(a) - L(d) + M(n) - M(e) - OL(m,d) - OM(m,e) - GL(m,d) - GM(m,e) = \Delta O + \Delta L + \Delta M - OL(m,d) - OM(m,e) - GL(m,d) - GM(m,e)$ (Moroccan man, disadvantaged location, media exposure gives ethnic effect compounded with location effect compounded with the media exposure effect)

The entire empirical model, including the set of identification constraints, can be written in matrix form as follows: $D = C*B$. D is an 11*1 column vector whose elements are the percentage point differences that are enumerated just above (i.e., the auxiliary parameters). C is an 11*9 matrix that imposes the constraints, and B the 9*1 vector whose elements are the parameters reflecting the discrepancies for which we seek direct estimates (i.e., the parameters of interest). More precisely:

$$D' = [\Delta(f,w,a,n), \Delta(m,w,a,n), \Delta(m,m,a,n), \Delta(f,w,d,n), \Delta(f,m,d,n), \Delta(m,w,d,n), \Delta(m,m,d,n), \Delta(f,w,d,e), \Delta(f,m,d,e), \Delta(m,w,d,e), \Delta(m,m,d,e)]$$

$$B' = [\Delta O, \Delta G, \Delta L, \Delta M, OG(m,w), OL(m,d), OM(m,e), GL(m,d), GM(m,e)]$$

$C =$

$$\begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & -1 & 0 & 0 & 0 & 0 \end{pmatrix}$$

$$\begin{array}{cccccccccc}
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & -1 & 0 \\
1 & 1 & 1 & 0 & -1 & -1 & 0 & 0 & 0 & 0 \\
1 & 0 & 1 & 0 & 0 & -1 & 0 & -1 & 0 & 0 \\
0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 1 & 0 & 0 & 0 & -1 & -1 & 0 \\
1 & 1 & 1 & 1 & -1 & -1 & -1 & 0 & 0 & 0 \\
1 & 0 & 1 & 1 & 0 & -1 & -1 & -1 & -1 & 0
\end{array}$$

Since the number of derived parameters (11) exceeds the number of structural parameters, the order condition for identification is fulfilled. With regards to the rank condition, C must be of rank 9 in order to ensure that C'C is of full rank. Given the actual values for the elements of C'C, computation using the SAS procedure IML gives a value of 40 for the det (C'C), which means that our rank condition is fulfilled.

The technique of asymptotic least squares estimation is applied to our model following an approach set out in Gouriéroux, Monfort and Trognon (1982, 1985) and Chamberlain (1982, 1984). More precisely, we estimate $\hat{D} = CB + u$, where $u = \hat{D} - D$ is derived by replacing the theoretical probabilities with their empirical counterparts. Let $\hat{V} = V(\hat{D})$ be the estimated variance-covariance matrix for \hat{D} . The optimal estimator of B, the vector of derived probabilities, is given by the feasible generalized least squares expression:

$$\hat{B} = (C' \hat{V}^{-1} C)^{-1} C' \hat{V}^{-1} \hat{D}, \text{ for which the variance-covariance matrix can be estimated by:}$$

$$\hat{V}(\hat{B}) = (C' \hat{V}^{-1} C)^{-1}.$$

The statistic for assessing the condition of over-identification can be computed from the empirical counterpart of u . Letting $\hat{u} = \hat{D} - C\hat{B}$, we can test for over-identification using the statistic $S = \hat{u}' \hat{V}^{-1} \hat{u}$, which has a chi-squared distribution with 2 (= 11 - 9) degrees of freedom. The results,

which are reported at the bottom of the following table, indicate that the over-identification tests have high P-values, which is interpreted as evidence that the model is appropriate. In practice, the estimator that we obtain generalizes the differencing estimators to the over-identified case, which we demonstrate in the appendix.

We first estimate the model including all nine of the discrepancy parameters (results presented in specification model 1, Table 3). In order to create additional constraints, we omit these parameters having the highest p-values and re-estimate to generate the results from specification model 2.

[Here table 3]

Turning to the results, we first discuss the estimates for the four direct, first-order discrepancy parameters. There is no evidence of discrimination with respect to ethnic origin or to location. On the other hand, the estimates for the pure gender effect (second row) and the pure media exposure effect (fourth row) are statistically significant. All other factors held constant, women are 4 percentage points *more* likely to receive a call-back, while those residing in the area with high media exposure are 3.2 percentage points less likely to receive a callback.

We now turn to the estimates for the interaction terms. The first two that are listed work to *increase* the degree of discrimination, as their sign is negative. The value of -0.033 listed in row five suggests that women of foreign origin experience a handicap such that the advantage that they hold by virtue of being a woman is mostly offset due to being of Moroccan origin. The value of - 0.040 listed in row six indicates that being of foreign origin and living in a disadvantaged area worsens discrimination by 4 percentage points.

On the other hand, the estimates of the remaining two interaction terms work to *decrease* the degree of discrimination, as their signs are positive. Being of Moroccan origin and residing in an area with unfavourable media exposure would reduce the degree of discrimination by 3.1 percentage points (row 8), but this estimate must be interpreted in conjunction with pure media exposure effect already mentioned. Living in a media-exposed area increases the degree of discrimination for everyone by 3.2 percentage points (the pure effect), but this effect is effectively offset for workers of Moroccan origin.

It appears as though the negative effect of unfavourable media coverage impacts primarily candidates of French origin. Men living in disadvantaged suburbs suffer less from discrimination to the tune of 3.8 percentage points (row 8), but this counter-intuitive discrepancy is interpreted in relation to the pure gender effect. Overall, men are less favored than women in the labor market (by -4 percentage points), but this discrepancy is compensated for by the effect of residing in an underprivileged area (3.8 percentage points). The upshot of those two estimates is that women lose their overall advantage over men when they live in disadvantaged areas.

We now mention some of the possible sources of bias in our estimates. In the context of labour market discrimination, an advantage of the technique of correspondence testing is that information on all of the candidates competing for a specific position is incorporated into the estimates, which should reduce the incidence of omitted variable bias that is inherent in non-experimental studies based on conventional data bases. On the other hand, inference is almost always based on a single cross-section of data on the outcome variable. In our context, it is possible that the media-exposure effect that we purport to capture is confounded with other observable influences between the towns of Sarcelles and Villiers-le-Bel. In many situations a difference-in-difference estimation strategy is executed in order to net out common trend effects and sharpen identification of the treatment effect. In our context, the ideal design would have been to conduct the same experiment before the riots occurred and perhaps for a third time a year after that event, but that was infeasible given the unpredictable and spontaneous nature of the event. In our view, the inherent strengths of the correspondence testing approach, the removal of the influences of heterogeneity between job postings through differencing operations, and the sharply identified difference estimator that we employ have addressed many but not all potential sources of bias.

6. Conclusion

We investigate a specific channel for effects of urban riots effects operating on the demand side of the labour market that would weigh heavily on the prospects of workers residing in affected areas, i.e. labour market discrimination based on locality. We assess the degree of potential discrimination using the “correspondence testing” technique with an experimental design. Our analysis is centered on the event of the riots that occurred in France in 2007, which were both intense and very geographically

concentrated in certain municipalities. The town of Villiers-le-Bel serves as our treatment unit, as it received a high degree of unfavourable media exposure. We selected two other nearby towns, Sarcelles and Enghien-les-Bains, to serve as control groups. Sarcelles is contiguous to Villiers-le-Bel, has a similar socio-economic-demographic profile, and was affected by the violence without attracting much media attention. Like Villiers-le-bel, it is considered to be disadvantaged, whereas Enghien-les-Bains is viewed as advantaged and was isolated from riot activity.

We discerned disparities in call-back rates over four dimensions: gender, ethnic origin, location, and the degree of media exposure during the riots. We decomposed the probability of receiving a callback for any candidate of given characteristics as the sum of several components, some of which reflect simple, one-way effects and others reflecting interacted effects. Those parameters measuring the degree of discrimination have been evaluated through simple difference estimators as well as difference-in-differences estimators computed by the estimation technique of asymptotic least squares.

Our findings confirm that riots can have significant consequences on labor market outcomes in the form of discrimination for workers having certain characteristics, and that these effects differ according to gender and ethnic origin. Specifically, we uncovered statistically significant negative effects related to pure media exposure effect. All other factors held constant, people residing in the area affected by riots and who were the subject of negative media coverage are 3.2 percentage points less likely to receive a callback. There is an empirical pattern underlying this discrepancy, namely that men of Moroccan origin, who incidentally are featured the most in the media coverage of the riots, are the least affected by the media-coverage effect once account is taken of interactions. On the other hand, it is women of French origin residing in that town that have been the most negatively affected by it. While the following statement is admittedly only a conjecture, a possible interpretation is that this pattern is consistent with a red-lining story based on town boundaries (as opposed to being based on pure distance). French women living in Villiers-le-Bel might be viewed by potential employers as particularly anomalous and hence risky candidates.

References

- Abadie, Alberto and Sofia Dermisi (2008) "Is terrorism eroding agglomeration economies in central business districts? Lessons from the office real estate market in downtown Chicago" *Journal of Urban Economics* 64(2), pp. 451-463.
- Arbel, Yuval and Ben-Shahar, Danny and Gabriel, Stuart A. and Tobol, Yossef (2010) "The Local Cost of Terror: Effects of the Second Palestinian Intifada on Jerusalem House Prices" *Regional Science and Urban Economics*, Vol. 40, pp. 415-426
- Arrow K.J. (1973) "The theory of discrimination", in O.A. Ashenfelter and A. Rees (eds.), *Discrimination in labor markets*, Princeton University Press, pp. 3-33.
- Becker G. (1957) *The economics of discrimination*, University of Chicago Press, (2nd edition: 1971).
- Bertrand, M. and Mullainathan, S. (2004) "Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination", *American Economic Review*, 94(4), pp. 991-1013.
- Besley, T., and Mueller, H.. (2012). "Estimating the Peace Dividend: The Impact of Violence on House Prices in Northern Ireland" *American Economic Review*, 102(2), pp. 810-33.
- Braakmann, Nils (2012). "The effect of the 2011 London riots on crime, policing and unemployment" MPRA paper No. 44883
- Bunel, M., L'Horty, Y. and Petit, P. (2016) "Discrimination based on place of residence and access to employment", *Urban Studies*, 53(2), pp. 267-286.
- Chamberlain, G. (1982) "Multivariate regression models for panel data", *Journal of Econometrics*, 18, pp. 5-46.
- Chamberlain, G. (1984). "Panel data" in Z. Griliches, and M. D. Intriligator (eds), *Handbook of Econometrics*, Vol. II, Elsevier, Amsterdam, pp. 1247-1318.
- Collins, W. and Margo, R. (2004) 'The labour market effects of the 1960s riots', in W. Gale and J. Pack, eds., *Brookings-Wharton Papers on Urban Affairs*, pp. 1-46, Washington DC: Brookings Institution.
- Collins, W. and Margo, R. (2007). "The economic aftermath of the 1960s riots in American cities: evidence from property values", *Journal of Economic History*, vol. 67(4), pp. 849-83.
- Détang-Dessendre, C. and Gaigné, C. (2009) "Unemployment duration, city size, and the tightness of the labor market", *Regional Science and Urban Economics*, 39, pp. 266-276.
- DiPasquale, Denise and Edward L. Glaeser (1998) "The Los Angeles Riots and the Economics of Urban Unrest" *Journal of Urban Economics* 43(1), pp. 52-78.
- Duguet E., Du Parquet, L. L'Horty, Y., Petit P. (2015). "New Evidence of Ethnic and Gender discriminations in the French Labor Market using experimental data: A ranking extension of responses from correspondence tests", *Annals of Economics and Statistics*, n° 117-118, pp 21-39.
- Duguet, E., Léandri, N., L'Horty, Y. and Petit, P. (2010) "Are Young French Job Seekers of Ethnic Immigrant Origin Discriminated Against? A Controlled Experiment in the Paris Area", *Annals of Economics and Statistics*, 99-100: 187-215.
- Fontaine F. (2006), «Les réseaux de relations : quelles perspectives pour l'économie du marché du travail ? », *Revue française d'économie*, 21 (1), pp. 127-172.
- Galster G. (2012), "The mechanism(s) of neighbourhood effects: theory, evidence, and policy implications", in M. van Ham, D. Manley, N. Bailey, L. Simpson and D. Maclennan (eds.), *Neighbourhood effects research: new perspectives*, Dordrecht, Springer, pp. 23-56.
- Gauschet, F. and Gaussier, N. (2004) "Urban Segregation and Labour Markets Within the Bordeaux Metropolitan Area: An Investigation of the Spatial Friction" Cahiers du GRES, N° 2004-19.
- Gautier P. and Zenou, Y. (2010), "Car ownership and the labor market of ethnic minorities", *Journal of Urban Economics*, 67 (3), pp. 392-403.
- Gobillon, L., Selod, H. and Zenou, Y. (2007) "The Mechanisms of Spatial Mismatch", *Urban Studies*, 44(12), pp. 2401-2427.

- Gobillon, L., Magnac, T. and Selod, H. (2011) “The effect of location on finding a job in the Paris region”, *Journal of Applied Econometrics*, 26(7), pp. 1079-1112.
- Gouriéroux, C., A. Monfort, and A. Trognon (1982). « Nonlinear asymptotic least squares ». INSEE, Document de travail no. 8207.
- Gouriéroux, C., A. Monfort, and A. Trognon (1985). « Moindres carrés asymptotiques ». *Annales de l'INSEE*, 58, pp. 91-122.
- Hellerstein J. and Neumark D. (2011), “Employment in black urban labor markets: problems and solutions” *NBER Working Papers* 16986, National Bureau of Economic Research.
- Hellerstein J., Kutzbach M. et Neumark D. (2014), “Do labor market networks have an important spatial dimension?“, *Journal of Urban Economics*, 79 (C), pp. 39-58.
- Huet-Vaughn E. (2015), “Quiet Riot: Estimating a Causal Effect of Protest Violence”, mimeo, Middlebury College
- Ioannides Y. et Datcher Loury L. (2004), “Job information networks, neighborhood effects, and inequality”, *Journal of Economic Literature*, 42 (4), pp. 1056-1093.
- L’Horty Y., Duguet E., Du Parquet L., Petit P. and Sari, F. (2011), « Les effets du lieu de résidence sur l’accès à l’emploi : un test de discrimination auprès de jeunes qualifiés », *Économie et Statistique*, 447, pp. 71-95.
- Matheson, Victor and Robert Baade (2004) “Race and riots: A note on the economic impact of the Rodney King riots”. *Urban Studies* 41(3), pp. 2691-2696.
- Neumark D., Bank R.J. et Van Nort K.D. (1996), “Sex discrimination in restaurant hiring: an audit study”, *Quarterly Journal of Economics*, 111 (3), pp. 915-941.
- Neumark, D. (2012). “Detecting Discrimination in Audit and Correspondence Studies” *Journal of Human Resources*. Vol. 47 Issue 4, pp. 1128-1157.
- Oreopoulos, P. (2011) “Why Do Skilled Immigrants Struggle in the Labor Market? A Field Experiment with Thirteen Thousand Resumes”, *American Economic Journal: Economic Policy*, 3(4): 148-71.
- Patacchini, E. and Zénou, Y. (2005) “Spatial mismatch, transport mode and search decisions in England”, *Journal of Urban Economics*, 58, pp. 62–90.
- Petit P., Duguet E., L’Horty Y., du Parquet, and Sari F. (2013). « Discrimination à l’embauche : les effets du genre et de l’origine se cumulent-ils systématiquement ? » *Economie et Statistique*, 464-465-466, pp. 141-153.
- Phelps E.S. (1972), “The statistical theory of racism and sexism” *American Economic Review*, 62 (4), pp. 659-661.
- Riach, P.A. and J. Rich, (2002) “Field experiments of discrimination in the market place”, *Economic Journal* 112, F480–F518.
- Tunstall, R., Green, A., Lupton, R., Watmough, S. and Bates, K. (2014) “Does Poor Neighbourhood Reputation Create a Neighbourhood Effect on Employment? The Results of a Field Experiment in the UK”, *Urban Studies*, 51(4), pp. 763–780.
- Zenou Y. (2002), “How do firms redline workers?”, *Journal of Urban Economics*, 52 (3), pp. 391-408.
- Zenou Y. (2006), Urban labor economic theory, in R. Arnott and D. McMillen (eds.), *A companion to urban economics*, Blackwell Publishing, chap. 25, pp. 418-439.

APPENDIX: An alternative, explicit approach to identification

This appendix shows another approach to deal with the over-identification issue that is inherent in our estimation procedure. Instead of computing the rank of the identification matrix, we derive an explicit way to get at the parameters of interest. This approach relies heavily on calculating differences in differences, which have an intuitive interpretation.

Gender discrimination. This particular quantity can be measured by comparing the call-back rates of men and women residing in the town of Enghien, which serves as our reference case, i.e. an advantaged town that did not experience unfavourable media coverage, netting out the effect of ethnic origin. Based on the assumption that there is no joint effect for the workers of French origin living in Enghien, we obtain the following estimator for the degree of **gender discrimination**:

$$(1) p(f,m,a,n) - p(f,w,a,n) = G(m) - G(w)$$

Origin discrimination. Enghien can also serve as a suitable benchmark for deriving the direct effect related to ethnic origin. A comparison of the male candidates (holding the other three factors constant) yields a measure of the degree of discrimination by ethnic origin.

$$(2) p(f,m,a,n) - p(m,m,a,n) = O(f) - O(m)$$

Origin-gender joint discrimination. Comparing the female candidates of both ethnic origins, while holding the other three factors constant, we obtain:

$$p(m,w,a,n) - p(f,w,a,n) = O(m) - O(f) + OG(m,w)$$

Replacing $O(m) - O(f)$ by the negative of its expression in equation (2) above and re-arranging, we obtain the following difference-in-differences estimator:

$$(3) OG(m,w) = p(m,w,a,n) - p(f,w,a,n) - [p(m,m,a,n) - p(f,m,a,n)] \\ = p(f,m,a,n) - p(f,w,a,n) - [p(m,m,a,n) - p(m,w,a,n)]$$

The joint effect of being a Moroccan-origin woman is estimated by comparing the difference of treatment between men and women of French origin with the difference of treatment between men and women of Moroccan origin.

Gender-Location joint discrimination. In order to assess this effect, we need to incorporate information about disadvantaged towns. We take the town of Sarcelles as the baseline case because it is considered to be disadvantaged, but it was not the subject of adverse media coverage. Comparing the French-origin candidates by gender, we first obtain:

$$p(f,m,d,n) - p(f,w,d,n) = G(m) - G(w) + GL(m,d)$$

Making a substitution from equation (1), we obtain the following expression:

$$(4) GL(m,d) = p(f,m,d,n) - p(f,w,d,n) - [p(f,m,a,n) - p(f,w,a,n)] ,$$

In this case we rely on the difference between the advantaged towns (on the right) and the disadvantaged towns (on the left) in the gender treatment differences.

Origin-Location joint discrimination. In this instance we still take the town of Sarcelles as the baseline case, but we now compare the two male candidates in order to incorporate the ethnic origin information:

$$p(m,m,d,n) - p(f,m,d,n) = O(m) - O(f) + OL(m,d)$$

Substituting from equation (2) and rearranging, we obtain:

$$(5) OL(m,u) = p(f,m,a,n) - p(m,m,a,n) - [p(f,m,d,n) - p(m,m,d,n)]$$

This expression consists of a comparison of the ethnic origin treatment difference between advantaged and disadvantaged towns.

Location discrimination: The estimation of this effect can be obtained simply by comparing the success rates of the French-origin female candidates, since we assume that they are not affected by joint discrimination effects. Comparing the outcomes from the towns of Sarcelles and Enghien, we obtain:

$$(6) p(f,w,a,n) - p(f,w,d,n) = L(a) - L(d)$$

Media discrimination: In a process that is symmetric to the derivation of equation (6), we compare the call back rates for French-origin female candidates between the disadvantaged towns of Sarcelles (no unfavorable coverage) and Villiers (unfavorable coverage).

$$(7) p(f,w,d,e) - p(f,w,d,n) = M(e) - M(n)$$

We can now obtain the two remaining joint effects, namely ethnic origin-media exposure and gender-media exposure.

Origin-media exposure joint discrimination: We compare the call-back rates for Moroccan-origin females between the towns of Villiers and Sarcelles:

$$(8) p(m,w,d,e) - p(m,w,d,n) = M(e) - M(n) + OM(m,e)$$

Substituting from equation (7) and rearranging, we obtain:

$$OM(m,e) = p(m,w,d,e) - p(m,w,d,n) - [p(f,w,d,e) - p(f,w,d,n)]$$

In this equation we compare the difference of treatment among the Moroccan-origin women between the two towns with the difference of treatment among French-origin women between the same two towns.

Gender-media discrimination: The final measure that we derive involves the interacted effects of gender and unfavourable media exposure. To this end we carry out two comparisons, the first of which involves Moroccan male candidates in the towns of Villiers (unfavourable exposure) and Sarcelles:

$$(9a) p(m,m,d,e) - p(m,m,d,n) = M(e) - M(n) + OM(m,e) + GM(m,e)$$

Substituting from equation (8) for the first three terms on the right-hand side, we obtain:

$GM(m,e) = p(m,m,d,e) - p(m,m,d,n) - [p(m,w,d,e) - p(m,w,d,n)]$. This expression gives the difference in treatment between Moroccan-origin males in Villiers and Sarcelles with the treatment difference between Moroccan-origin females in the same two towns.

Now consider the corresponding effect for French-origin men:

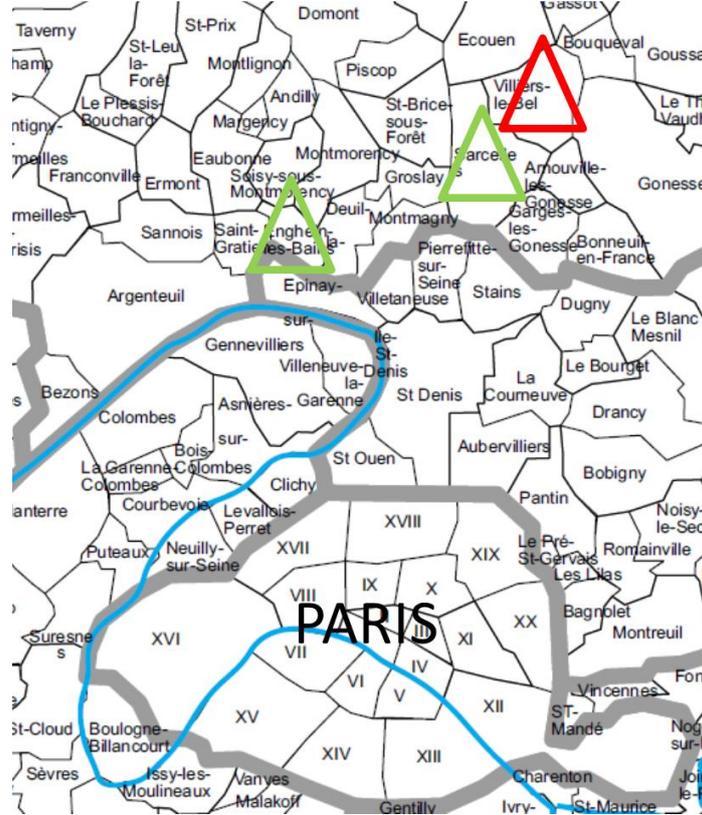
$$p(f,m,d,e) - p(f,m,d,n) = M(e) - M(n) + GM(m,e)$$

Substituting from equation (7), the counterpart for equation (9a) is expressed as:

$$(9b) GM(m,e) = p(f,m,d,e) - p(f,m,d,n) - [p(f,w,d,e) - p(f,w,d,n)]$$

This equation creates an over-identification restraint.

Map 1. Geographical Location of the treatment area (Villiers-le-Bel) and the control areas (Sarcelles and Enghien-les-bains)



Source: Administrative map of the Paris region

Figure 1. Times series data on the number of “Google hits” for the three areas

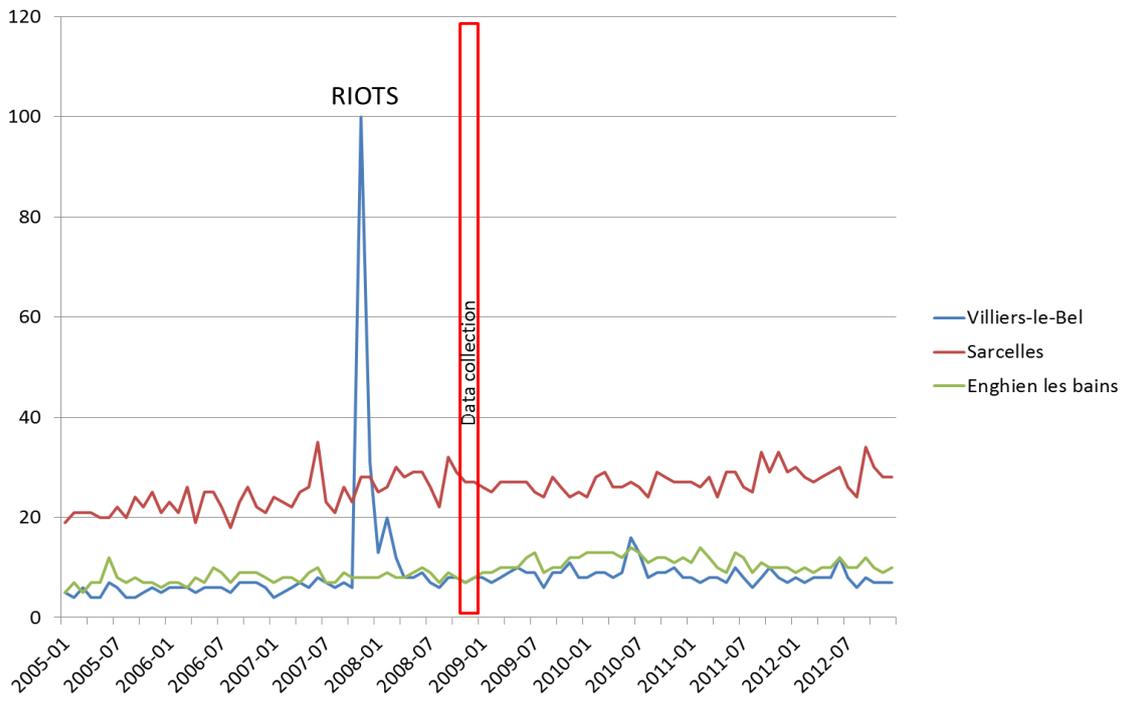


Table 1 : Call back rates

	Favourable reponse rate	t-statistic	90 % confidence interval	
			Lower bound	Upper bound
French Origin				
<i>women :</i>				
Enghien	22,5% **	9,39	18,6%	26,4%
Sarcelles	22,1% **	9,31	18,2%	26,1%
Villiers-le-Bel	17,9% **	8,23	14,3%	21,5%
<i>men :</i>				
Enghien	16,9% **	7,95	13,4%	20,5%
Sarcelles	19,9% **	8,76	16,3%	23,8%
Villiers-le-Bel	18,6% **	8,40	15,0%	22,1%
Moroccan Origin				
<i>women:</i>				
Enghien	19,5% **	8,71	16,0%	23,1%
Sarcelles	13,7% **	6,96	10,4%	16,9%
Villiers-le-Bel	15,0% **	7,31	11,7%	18,2%
<i>men :</i>				
Enghien	18,6% **	8,45	15,0%	22,1%
Sarcelles	19,2% **	8,58	15,6%	23,1%
Villiers-le-Bel	17,3% **	7,96	13,7%	20,8%
Global positive response rate¹	52.1%			

1. Percentage of offers for which fictitious candidates received at least positive response

The t-statistics and the confidence intervals were calculated by the bootstrapping method executed over 10,000 trials; * significant at the 10% level; ** significant at the 5 % level.

Table 2: Probabilities components

Candidate	Origin	Gender	Location	Media	Origin* gender	Origin* location	Origin*me dia	Gender* location	Gender* media
Town of Enghien									
French, woman P (f,w,a,n)	O(f)	G(w)	L(a)	M(n)					
French, man P(f,m,a,n)	O (f)	G(m)	L(a)	M(n)					
Moroccan, woman P (m,w,a,n)	O(m)	G(w)	L(a)	M(n)	OG(m,w)				
Moroccan, man P (m,m,a,n)	O(m)	G(m)	L(a)	M(n)					
Town of Sarcelles									
French, woman P (f,w,d,n)	O(f)	G(w)	L(d)	M(n)					
French, man P (f,m,d,n)	O(f)	G(m)	L(d)	M(n)				GL(m,d)	
Moroccan, woman P (m,w,d,n)	O(m)	G(w)	L(d)	M(n)	OG(m,w)	OL(m,d)			
Moroccan, man P (m,m,d,n)	O(m)	G(m)	L(d)	M(n)		OL(m,d)		GL(m,d)	
Town of Villiers-le-Bel									
French, female P (f,w,d,e)	O(f)	G(w)	L(d)	M(e)					
French, male P (f,m,d,e)	O(f)	G(m)	L(d)	M(e)				GL(m,d)	GM(m,e)
Moroccan, female P (m,w,d,e)	O(m)	G(w)	L(d)	M(e)	OG(m,w)	OL(m,d)	OM(m,e)		
Moroccan, man P(m,m,d,e)	O(m)	G(m)	L(d)	M(e)		OL(m,d)	OM(m,e)	GL(m,d)	GM(m,e)

Table 3: Estimation of the callback probabilities components

Parameters	Estimation	Standard error	P value	Estimation	Standard error	P value
Model		(1)			(2)	
ΔO	-0.011	0.022	0.614			
ΔG	-0.051	0.022	0.019	-0.040	0.019	0.035
ΔL	0.010	0.021	0.641			
ΔM	0.031	0.019	0.109	0.032	0.015	0.033
OG (m,f)	-0.042	0.022	0.060	-0.033	0.017	0.059
OL (m,u)	-0.040	0.024	0.088	-0.040	0.016	0.012
OM (m,e)	0.026	0.022	0.233	0.031	0.020	0.127
GL (m,u)	0.045	0.026	0.083	0.038	0.022	0.079
GM (m,e)	0.005	0.022	0.836			
Constraints		2			5	
S		1.73			2.52	
P value		0.42			0.74	