

Prosecution or Persecution? Extraneous Events and Prosecutorial Decisions

Non Peer-reviewed author version

BIELEN, Samantha & Grajzl, Peter (2021) Prosecution or Persecution? Extraneous Events and Prosecutorial Decisions. In: Journal of empirical legal studies (Print), 18(4), p. 765-800..

DOI: 10.1111/jels.12301

Handle: <http://hdl.handle.net/1942/35516>

Prosecution or Persecution? Extraneous Events and Prosecutorial Decisions*

Samantha Bielen
Faculty of Business Economics
Hasselt University
3500 Hasselt, Belgium
Email: samantha.bielen@uhasselt.be

Peter Grajzl (corresponding author)
Department of Economics
The Williams School of Commerce, Economics and Politics
Washington and Lee University
204 West Washington St.
Lexington, VA 24450, USA
and
CESifo
Munich, Germany
Phone: +1 540 458 8049
Fax: +1 540 458 8639
Email: grajzlp@wlu.edu

* For helpful comments and suggestions, we thank Naci Mocan, participants at the Hamburg Lectures in Law & Economics seminar and the Judicial Politics Workshop, and two anonymous reviewers.

Prosecution or Persecution? Extraneous Events and Prosecutorial Decisions

August 16, 2021

Abstract

In 2004 a Dutch-Moroccan Islamic extremist in Amsterdam brutally murdered Theo van Gogh, a Dutch filmmaker known for his critique of Islam. Using a difference-in-difference approach, we show that, in Amsterdam, the assassination resulted in more than a 19-percentage-point increase in the likelihood of prosecution for unrelated violent crimes of male suspects born in Muslim-majority countries. The effect is detectable during the first month following the murder but dissipates thereafter. We find no evidence of the murder's effect for non-violent crimes and for violent crimes processed by other Dutch prosecution offices. Our findings are thus not consistent with a purely emotions-based explanation of the murder's effect. Instead, our results are congruent with both a signaling explanation, whereby career-motivated prosecutors chose to showcase their toughness, and a behavioral explanation entailing prosecutors' susceptibility to availability heuristic bias. On institutional grounds, we view the latter as more plausible than the former. Our paper adds to an emerging literature demonstrating that extraneous events can critically shape criminal justice outcomes.

Keywords: criminal prosecution, extraneous events, ethnoreligious bias, justice, the Netherlands

JEL Classifications: K40, K14, J15, D91, P48

1. Introduction

That extraneous events and emotional shocks can influence human decisions has been well acknowledged in the literature on social psychology. Salient events, such as rare accidents or instances of extreme violence, raise individuals' perceived frequency of comparable incidents. Easily recalled, salient events thereby shape individuals' assessments and choices in subsequent separate contexts, a manifestation of the availability heuristic bias (Tversky and Kahneman 1973, Stalans 1993). Similarly, events that trigger negative emotions, such as sadness and anger, have been found to impact judgment not only with respect to the specific circumstances that provoked the initial emotional response, but, importantly, also in entirely unrelated settings (Lerner et al. 2015).

In this paper we demonstrate that decision-makers' responses to salient extraneous events can be an important driver of behavior also in the context of criminal prosecution, a crucial but comparatively understudied stage of the criminal justice process where information irrelevant to the case at hand and decision-makers' personal feelings should, by law and on equity grounds, decidedly not influence the outcomes. The extraneous event that is the focus of our analysis is the 2004 killing of Theo van Gogh, a Dutch filmmaker known for his critical and provocative stance toward Islam. Van Gogh was brutally murdered on November 2, 2004 in Amsterdam by Mohammed Bouyeri, a Dutch-Moroccan Islamic extremist. The assassination caused a public outrage, led to violent retaliations against the Dutch Muslims, and added a new layer of aggression to the already heated debate about multiculturalism and the role of Islam in the Netherlands. The Amsterdam government, in particular, found itself under siege for its allegedly overly tolerant and inclusive stance toward minorities.

Utilizing administrative data on the universe of all criminal charges processed by the Amsterdam prosecution office around the time of the van Gogh murder and employing a difference-in-difference approach, we show that, following van Gogh's assassination, the prospects of prosecution for unrelated violent crimes of male suspects born in Muslim-majority countries increased by more than 19 percentage points. The corresponding effect is detectable for the first month after the van Gogh murder, but dissipates thereafter. Furthermore, the effect is discernible even when we restrict our attention only to charges stemming from crimes committed before the van Gogh murder, an approach immune to concerns that the killing perhaps changed in some

unobserved way the nature of the underlying crimes. Placebo tests and the application of Oster's (2019) methodology show that our results cannot be attributed to omitted variables.

Interestingly, we find no comparable effect for non-violent crimes and for violent crimes processed by other, non-Amsterdam prosecution offices in the Netherlands. These results cannot be explained by prosecutors' negative emotional reactions. Had prosecutorial decisions been driven by feelings of anger against male Muslims, we should have detected the effect for multiple crime types and likely beyond Amsterdam. Instead, our findings are more congruent with either a signaling or a behavioral explanation entailing the availability heuristic bias (Tversky and Kahneman 1973). According to the former (signaling) explanation, career-motivated prosecutors showcased their toughness to their revenge-motivated political principals or perhaps directly to the male Muslim minority at large. According to the latter (behavioral) explanation, the murder induced an association in prosecutors' minds between male Muslim suspects and culpability for violent crimes. By consciously making use of the corresponding perceptual shorthand, the prosecutors were subsequently more likely to press charges against the pertinent suspects. Our data and empirical results do not allow us to distinguish between these competing mechanisms. We suggest that, on institutional grounds, the availability heuristic-based explanation appears perhaps more plausible than the signaling explanations.

Our paper thereby contributes, first and foremost, to the growing empirical literature on the influence of extraneous events and emotions for administration of justice. Within this literature, attention has nearly exclusively focused on either judicial or jury decision-making, both in criminal and in civil context (e.g., Eren and Mocan 2018, Philippe and Ouss 2018, Brodeur and Wright 2019, Danziger et al. 2011, McConnell and Rasul 2020, Heyes and Saberian 2019, Shayo and Zussman 2011, Philippe and Ouss 2018).¹ In contrast, the role of extraneous incidents for prosecutorial decision-making has received very limited consideration. McCannon and Wilson (2019) show that wrongful rape charges against one U.S. university's lacrosse team affected the propensity to prosecute rape of state prosecutors facing elections. Rasul and McConnell (2020) find that, in the U.S., the anti-Muslim sentiment induced by 9/11 terrorist attacks impacted prosecutors' attitudes towards Hispanics. As one of the first contributions to reveal the importance

¹ A related criminal-justice literature examines the presence of extralegal disparities with respect to defendant and judge characteristics, focusing especially on race and gender. See, e.g., Cohen and Yang (2019), Depew et al. (2017), Starr (2015), Lim et al. (2016), and Fischman and Schanzenbach (2012).

of extraneous events for prosecutor behavior, our paper shows that Bouyeri's murder of van Gogh in Amsterdam temporarily increased the likelihood of prosecution for unrelated violent crimes of male suspects born in Muslim-majority countries.

Investigation of prosecutorial decisions is productive not only on its own substantive merits, but also because it offers a distinct methodological advantage over the study of judicial decisions. Empirical research of judicial behavior in the criminal-justice context is often plagued by concerns that the sample of tried cases is not a random sample of all recorded criminal cases. Thus, effects attributed to judges and juries based on examination of court data alone need not always be a consequence of judge and jury behavior, respectively. Rather, they can be a reflection of decisions made in the earlier, prosecution stage of the criminal justice process (see, e.g., Rehavi and Starr 2014, Rasul and McConnell 2020). The focus on prosecutorial decision-making per se avoids these specific concerns. Furthermore, as we explain below, in the Netherlands the procedural rules for processing of violent criminal cases ensure that Dutch prosecutors have a monopoly over prosecution decisions. The police are obliged to hand over all violent criminal cases to the applicable prosecution office. Consequently, our focal sample is, for institutional reasons, also immune to a further layer of selection concerns that can arise in settings where the sample of charges processed by prosecutors is itself an outcome of law enforcement's earlier, deliberate processing of alleged crimes.

At the same time, our paper adds to the broader empirical literature on prosecutor behavior. The latter remains comparatively underexplored relative to the behavior of other actors in the criminal justice process (see, e.g., Garoupa 2012, Tonry 2013, Voigt and Wulf 2019). The overwhelming majority of the existing empirical literature on prosecutors has focused on the U.S. context. Because most U.S. state prosecutors are elected, not appointed, political considerations and election-based incentives have naturally been one very important focus in this strand of research (e.g., McCannon 2013, Bandyopadhyay and McCannon 2014, McCannon and Wilson 2019, Dyke 2007, Nadel et al. 2017).² In contrast, research on prosecutorial decision-making outside of the U.S. is comparatively scarce (see, e.g., Luna and Wade 2012, Gilliéron 2014). We contribute to this literature by examining the effect of one especially salient extraneous event for

² Other empirical research on federal and state prosecutors in the U.S. has focused, for example, on prosecutors' career outcomes and concerns (Boylan 2004, Boylan and Long 2005), incentives to dampen the effects of mandatory minimum sentencing laws (Bjerk 2005), and the allocation of cases to the state versus the federal systems (Glaeser et al. 2000).

prosecutorial treatment of minority suspects in the Netherlands, a jurisdiction where prosecutors are appointed civil servants.

The rest of the paper is organized as follows. Section 2 provides the necessary institutional and conceptual backgrounds. Section 3 introduces our data. Section 4 develops our empirical strategy and presents the main results. Section 5 offers a series of robustness checks and additional results. Section 6 concludes.

2. Institutional and Conceptual Background

2.1. The Theo van Gogh Murder and Its Aftermath

Theo van Gogh was a filmmaker, columnist and television host known for openly criticizing Islam and often distastefully provoking Muslims. Van Gogh challenged the existing political correctness and, in his view, the resulting "excessive" respect towards Muslim minorities and the Islam in particular (Hajer and Uitermark 2008). His inflammatory and often vulgar statements, which he defended on the grounds of freedom of speech, led to several slander and blasphemy lawsuits. In his short-film *Submission*, van Gogh drew on the Koran and a script written by Ayaan Hirsi Ali, a Somali-born Dutch member of parliament and activist who had rejected the Islamic faith, to criticize the treatment of women under Islam. After *Submission* aired on television in August 2004, both van Gogh and Hirsi Ali received death threats.

On November 2, 2004, van Gogh was murdered in Amsterdam by Mohammed Bouyeri, a second-generation Moroccan immigrant and member of a radical Islamist group the Hofstad Network. Bouyeri shot van Gogh eight times, cut his throat in an attempt to decapitate the victim, and stabbed a letter to van Gogh's dead body, threatening Hirsi Ali, two other politicians (including the mayor of Amsterdam), and Western governments in general. In a separate letter that Bouyeri attached to himself, he explained his religious motivations, including his desire for martyr's death. The police nevertheless managed to arrest him. In July 2005, Bouyeri was sentenced to life in prison.

The van Gogh murder triggered many hot-blooded reactions against Islam and Muslim minorities (see, e.g., Gautier et al. 2009, Van Atteveldt et al. 2005, De Haan 2007, Hajer and Uitermark 2008). Opponents of multiculturalism in the Netherlands quickly connected the event to the government's allegedly soft integration policies. The criticism was directed towards the national government, but many critics targeted the Amsterdam government specifically. The latter had acquired a reputation for its especially favorable stance towards multiculturalism and tolerant

leanings toward ethnoreligious minorities (Hajer and Uitermark 2008). Amidst nation-wide panic and anger, the overwhelming majority of political leaders called for moderation, referring to the fact that the Muslim minority itself had nothing to do with the murder, nor did they have the slightest sympathy for it (De Haan 2007). On the other hand, reactions of isolated individuals were aimed to further provoke the growing tensions. The Dutch minister of finance, for example, "referred to a clash of civilizations, declared that there was a war going on between Islam and the West, and boasted that if extremist Muslims wanted war, they could get it" (Van der Veer 2006). In the days after the murder, multiple mosques and Islamist schools were burnt down and tainted with anti-Muslim slurs, resulting in fears of widespread retaliation against Muslims (Hajer and Uitermark 2008, Van der Veer 2006).

The van Gogh murder was the key event that drove the public discourse on Islam and the Muslim minority in the Netherlands for several weeks. To illustrate this point, Figure 1 plots two time series capturing the daily incidence of accounts (articles, editorials, and other features) in eight major Dutch newspapers shortly before and after the murder.³ The first series, in black, captures the time evolution of the daily frequency of accounts containing words 'Islam' or 'Muslim' (and their derivative grammatical forms in Dutch) that also discuss van Gogh. The series entails an easily-discernible jump after November 2, 2004, an indication of the media's attention devoted to the van Gogh murder in connection with discussion of Islam and Muslim minorities. The downward trend in the series following the murder shows that the corresponding topic was very prominent in the news for several weeks until it eventually faded away. The second series, in gray, shows the time evolution of the daily frequency of accounts containing words 'Islam' or 'Muslim' that do not include 'van Gogh'. The corresponding series exhibits no clear trend and, importantly, features no structural break around November 2, 2004. This is an indication that on or around November 2, 2004, the van Gogh murder was the salient event that was especially pertinent to the societal discourse on Islam and the Muslim minority in the Netherlands.

Much like September 11, 2001 in the U.S., the van Gogh murder is believed to have changed the perceptions of Muslim minorities in the Netherlands. Surveys demonstrate that 80 percent of the respondents believed that the murder affected the relationship between Muslims and non-Muslims in the Netherlands (Van der Veer 2006). Prior research shows that the van Gogh

³ The newspapers are: Telegraaf, Algemeen Dagblad, Volkskrant, NRC Handelsblad, Trouw, Financieele Dagblad, Reformatorisch Dagblad, and Nederlands Dagblad.

murder led to a fall in housing prices in the Dutch neighborhoods in which more than a quarter of the residents originate from a Muslim country (Gautier et al. 2009). In several European countries, van Gogh's assassination was associated with a subsequent increase in citizen support for restrictive immigration policies (Finseraas et al. 2011). Yet to date, no study has explored the impact of the van Gogh murder on the administration of criminal justice, and in particular the prosecutorial decisions with regard to Muslim-minority suspects. To set the stage for the analysis, we first clarify the key elements of the prosecution process in the Netherlands and explain the possible behavioral mechanisms through which the van Gogh's killing could conceivably impact prosecutorial decisions vis-à-vis Muslim-minority suspects.

2.2. Prosecutors and the Prosecution Process

The Dutch public prosecution service operates under the auspices of the Minister of Justice and is governed by the Board of Prosecutors General.⁴ The core of the public prosecution service is organized into ten district-level public prosecutor offices with corresponding geographic jurisdiction. Each office is headed by a Chief Prosecutor and consists of specialized units based on crime type (e.g. organized, commercial, environmental). In addition, the public prosecution service has specialized offices for prosecution of particular types of crimes (e.g. organized crimes that require a coordinated approach). In this paper, we without loss of generality focus on the crimes processed by the ten main public prosecution offices.

To become a public prosecutor, a candidate must be a Dutch national, have a law degree, possess six years of relevant legal experience, and complete a training program. Unlike in the U.S., where state prosecutors are elected, Dutch prosecutors are appointed (by the Crown or the Minister of Justice). Dutch prosecutors are thus comparatively less exposed to public pressure to pursue charges and secure convictions. However, the prosecutors are subject to career concerns. Dutch prosecutors are public servants. Unlike judges, they are not appointed with life tenure. Prosecutors' remuneration is based on experience and, conditional on years of experience, regulated by a pay scale.⁵ Excellent performance is rewarded with a raise. If a prosecutor has achieved maximum pay within the applicable pay scale, further salary increases are possible via promotion to a higher rank. The latter can occur after a comprehensive performance review that culminates in an official report. Promotion is formally awarded by the minister of justice. If a prosecutor engages in

⁴ This section draws on Tak's (2008) overview of the criminal prosecution system in the Netherlands.

⁵ Legal Status of Judicial Officers Act (*Wet rechtspositie rechterlijke ambtenaren*).

misconduct or fails to act in ways expected from a public prosecutor, the prosecutor may be subject to disciplinary measures that include a reprimand, a fine, lower salary or no automatic raise, a suspension, or loss of employment. Depending on the severity of the transgression, disciplinary measures are imposed either by the pertinent prosecutorial office or the minister of justice. In sum, higher-level administrative and political principals are, at least in principle, able to wield some level of influence over the prosecutors.

The primary objective of the public prosecution service is criminal enforcement and deterrence. If a crime is committed, the police get involved either on the basis of a report or because an offender was caught red-handed. Only in the latter case, as well as when the reported crime carries a prison sentence of at least four years, can the police proceed to an arrest autonomously. Such an arrest can last up to six hours, and is intended to facilitate the interrogation of the suspect about the alleged crime.⁶ Importantly, regardless of whether and if so how a suspect is arrested by the police, a criminal case is always registered with the applicable public prosecutor office for further processing. That is, except in instances of minor wrongdoings (e.g. lesser traffic offences), the Dutch police do not have the legal mandate to drop a case without involvement of the prosecution office. In particular when it comes to all serious crimes, the police in the Netherlands are legally obliged to register the corresponding crimes with the prosecution office.

A criminal case registered with a prosecution office is allocated to a specific prosecutor based on the type of alleged crime, the caseload faced by individual prosecutors, and the prosecutors' pre-determined work-hours schedule. For example, a case involving commercial crime is allocated to the pertinent specialized unit, and within the unit to the specific team member available at the time of assignment. Thus, conditional on the nature of the crime under investigation, the assignment of cases among prosecutors within a given prosecutorial office should be on institutional grounds de facto independent of the characteristics of the prosecutor assigned the case.

Dutch public prosecutors have a monopoly and full autonomy over prosecutorial decisions. For a given charge in a specific criminal case involving a particular suspect, prosecutors can make one of three decisions: pursue the charge or charges, in which case the pertinent criminal case

⁶ In the case of certain crime types, a suspect can also be held in police custody for a maximum of three days, by order of the public prosecutor or a senior police officer. When the police-custody expires, the suspect can be either released or ordered to remand in custody for fourteen days. Such a custody is requested by a public prosecutor.

proceeds to court; drop a charge or charges (e.g. due to inadmissible evidence, for expediency reasons, or the passing of the suspect); or dispose the case without either pursuing or altogether dropping charges. In practice, the last (third) scenario encompasses both conditional suspension of prosecution and formally negotiated settlement referred to as the transaction (*transactie*). The latter mode of disposition is legally a possibility for the majority of processed crimes (those that carry a prison sentence of less than six years). Under transaction, the suspect agrees to pay a sum of money to avoid prosecution. In exchange, the crime is not reported on the suspect's criminal record. Note, however, that the transaction is not the Dutch counterpart of the American plea bargaining: a completed transaction does not imply that the offender has pleaded guilty to committing the crime.

If the prosecution office decides to pursue prosecution, the prosecutor drafts a summons that lists the specific charges for which the suspect is prosecuted and describes the constituent elements of the alleged punishable conduct. The issuing of the summons indicates the start of the trial stage of the criminal justice process.

Dutch prosecutors are expected to ensure impartiality and fairness of outcomes. Importantly, prosecutors are not mandated to motivate their decision if they choose not to press charges. In rendering decisions, Dutch prosecutors thus enjoy a considerable scope for discretion, a feature of the prosecution process that opens the door for the influence of extralegal factors.

2.3. How Might Have the van Gogh Murder Affected Prosecutorial Decision-Making?

There are multiple reasons why the assassination of Theo van Gogh could have impacted prosecutorial treatment of male Muslim-minority suspects irrespective of the underlying case facts. We first discuss two behavioral channels. We then consider a third set of mechanisms that stress political pressures and signaling.

One, inherently behavioral, mechanism stresses the importance of availability heuristic (Tversky and Kahneman 1973), a well-known behavioral bias detected in a variety of settings, including the administration of justice where decision-making normally follows a deliberative process (see, e.g., Stalans 1993, Diamond and Stalans 1989, Korobkin and Ulen 2000, Burns 2016, Brodeur and Wright 2019). According to this argument, Bouyeri's ruthless murder increased the prosecutors' perceived rate of incidence of violent crimes committed by Muslim immigrants, thereby eliciting an association in prosecutors' minds between male suspects from Muslim-majority countries and culpability for violent crimes. To the extent that prosecutors, faced with considerable costs of information-processing, rely on perceptual shorthands and heuristics when

forming attributions of guilt, the likelihood of prosecution of suspects from Muslim-majority countries subsequently increased. Importantly, the availability heuristic-based mechanism entails a combination of unconscious and conscious mental processes: while the murder-induced increase in the perceived frequency of culpability of Muslim males for violent crimes is unconscious, the decision to act on the heuristic, or to avoid revisiting one's perceptions, is a conscious one (Brodeur and Wright 2019: 122).

Another, also behavioral, channel through which van Gogh's killing could have influenced prosecutorial decisions emphasizes emotions. Research in psychology has shown that "incidental anger triggered in one situation automatically elicits a motive to blame individuals in other situations even though the targets of such anger have nothing to do with the source of the anger...Moreover, carryover of incidental emotions typically occurs without awareness" (Lerner et al. 2015: 803). For example, unexpected losses of an adjudicator's favorite college football team can result in more severe sentencing of juvenile offenders (Eren and Mocan 2018). Based on this theory, Bouyeri's killing of van Gogh caused an emotional shock, resulting in anger and resentment against the Muslim community. The shock spurred an anti-Muslim response not only among the general public, but also among educated public servants such as prosecutors who, whether consciously aware of their emotional state or not, ended up prosecuting male suspects from Muslim-majority countries more aggressively.

A third set of reasons why the murder could have resulted in harsher prosecution of male suspects from Muslim-majority countries is based on signaling considerations. Specifically, the prosecutors' more aggressive charging of male Muslim suspects after the murder could be a consequence of external pressures emanating either from the outraged public or, alternatively, the government administration as the prosecutors' political principal. That the former (the public at large) had a decisive sway is not likely: Dutch prosecutors are appointed, not elected, and thus not under direct public pressure to prosecute high-profile crimes. In contrast, the possibility that political principals exerted influence on prosecutorial behavior cannot be readily discarded.

Certainly, the legally-ensured autonomy at least nominally protects the prosecutors from politically-motivated interference.⁷ Yet at the same time, as public servants, Dutch prosecutors are

⁷ Similarly, the observed prosecutors' decisions should not be a mere unbiased reflection of anticipatory effects concerning subsequent biased judicial decisions. In particular, if following the van Gogh killing prosecutors expected judges to be biased against suspects from Muslim-majority countries, then unbiased prosecutors aiming to correct the anticipated judicial bias would have arguably chosen to prosecute those suspects less likely. As we clarify below, this

subject to career concerns (see Section 2.2) and hence not immune to top-down pressures stemming from the political-administrative hierarchy. To be sure, we are aware of no media or other narratives suggesting that, following the van Gogh killing, the core of the national or local political leadership in any way aimed to crack down on male Muslim criminal suspects. To the contrary, the predominant political rhetoric was decidedly one of de-escalation (see Section 2.1). But public rhetoric sometimes hides private truths (Kuran 1995). If some of the key administrative and political actors privately sought revenge, then the more aggressive charging of male Muslim suspects after the murder would have served as a vehicle for career-motivated prosecutors to signal their toughness and resolve directly to those principals. Alternatively, the prosecutors, possibly subject to an availability heuristic bias or under an emotional shock, may have chosen to display their power to the Muslim community at large irrespective of the wishes of their political principals. In either case, the decision to prosecute male Muslim suspects of violent crimes relatively more aggressively after the murder is consistent with a signaling explanation.

Of course, different prosecutors may differ in their susceptibility to above-mentioned mechanisms. As we clarify below, our data do not allow us to link specific cases to particular prosecutors and we do not observe the deliberative context of and justifications for prosecutorial decisions. But the data do allow us to examine the effect for a multiplicity of crime types and for all prosecution offices in the Netherlands. This allows us to probe the validity of at least some one of the above-discussed mechanisms. Specifically, a finding that the van Gogh murder affected prosecutorial treatment of the focal suspects only in Amsterdam, or only with respect to violent crimes, would not resonate with a pure emotions-based explanation. Had the van Gogh killing, which echoed strongly across all of the Netherlands (see Section 2.1), among prosecutors truly elicited especially powerful feelings of anger or resentment toward Muslim men, then we would have expected to observe a harsher treatment of the corresponding suspects across a range of crime types and multiple prosecution offices in the Netherlands. On the other hand, evidence that male suspects from Muslim-majority were post murder treated more harshly only with respect to violent crimes and only in Amsterdam is fully consistent with both the availability heuristic-based behavioral explanation and a signaling story. Because the violent murder of van Gogh took place in the city of Amsterdam, the scope for availability heuristic-based actions and any retaliation

is not what we observe in the data. On the other hand, prosecutors whose treatment of suspects from Muslim-majority countries had been influenced by the murder would have certainly opted to prosecute cases involving such suspects.

motives of the prosecutors' political-administrative principals would have naturally been the greatest in Amsterdam, especially soon after the murder. In sum, our data offers an opportunity to contrast the relevance of theories stressing the role of emotions versus theories emphasizing availability heuristic bias and signaling, but we are unable to credibly distinguish between the availability heuristic-based and the signaling explanations.

3. Data

3.1. Source and Sample

The source of our data is a (non-public) database maintained by Statistics Netherlands. The full database contains information on the population of criminal cases registered at and processed by the prosecution offices in the Netherlands between 2001 and 2017. As we clarify below, in our primary analysis we restrict our attention to violent criminal cases and associated charges that were processed by the Amsterdam prosecution office during the 12-week interval around the day of van Gogh's assassination (November 2, 2004) and that involve a male suspect. (In various robustness checks we use alternative samples.) We exclude all charges based on crimes committed on the day of the van Gogh murder (and thus charges stemming from van Gogh's murder itself). Upon dropping observations with missing values, our focal sample consists of 1,193 charge-level observations involving 766 suspects, of which 745 are distinct suspects.⁸

Importantly, our data entails the universe of all pertinent charges brought to and processed by the prosecution offices during the time period under consideration. Our data are therefore immune to sample selection concerns that invariably arise in studies of adjudicatory decisions made at the conviction and sentencing stages of the criminal justice process: cases processed by courts are not a random sample of all cases processed by prosecutors. In addition, because in the Netherlands the prosecution offices have a monopoly over prosecutions, our data are not subject to selection concerns that could arise if the police were involved in substantive case-processing decisions. In the Netherlands, the police have the authority to dismiss a case without involvement of the pertinent prosecution office only in a very limited set of minor drug cases. The police have no discretionary authority when it comes to prosecution of violent crimes (see Section 2.2).

Of course, selection issues can still arise if the police are more likely to register crimes involving specific suspects or if, as a consequence of ethno-religious profiling, law enforcement

⁸ Thus, a small set of suspects in our dataset is charged more than once, each time in relation to a different crime.

is comparatively more likely to detect crimes committed by targeted minorities. In the Netherlands, this is certainly a possibility with regard to crimes without a clearly identifiable victim (e.g. a purchase of drugs or prostitution): such crimes are normally not reported and hence the police play an active role in discovering them on the spot. However, violent crimes, which constitute the focus of our analysis, are overwhelmingly reported by the victims or witnesses. Consequently, the Dutch police have virtually no control over which specific violent crimes get registered and thus processed by the prosecution offices. That is, the sample of violent crimes that we work with should not be skewed by discretionary decisions that the police in some jurisdictions can make in the earliest stages of the law-enforcement process.

3.2. Key Variables and Descriptive Statistics

Table 1 provides definitions of the main variables. (In the estimation, we further employ multiple sets of fixed effects.) Table 2 reports charge-level descriptive statistics. For each charge, we know, first and foremost, the identity of the suspect and the prosecution decision of the applicable prosecution office. We define the suspect's status as non-native from a Muslim-majority country (the treated group) based on the suspect's recorded country of birth. In our main specification, we code a suspect as a non-native from a Muslim-majority country (a Muslim minority, in short) when the suspect is born in a country where at least 50 percent of the population identifies as Muslim (we examine the robustness of our findings to varying this threshold and using alternative definitions).⁹ From the charges in our primary sample, 17.7 percent involve Muslim-minority suspects. From the 745 distinct suspects in our primary sample, 139 were born in Muslim-majority countries. Our primary outcome of interest is whether the applicable prosecution office chose to prosecute the given charge as opposed to drop it or pursue alternative modes of disposition. 79 percent of charges in our primary sample are prosecuted.

We further know the suspect's age, employment status (unemployed, retired, a student or employed), and criminal history at the time of the crime. Measured at the charge level, average suspects age is 34.2 years. 57 percent of the suspects are unemployed. 43 of suspects possess a prior criminal history, defined as at least one conviction in the last three years. We have

⁹ Based on this threshold, the list of Muslim-majority countries that the suspects in our data hail from is: Afghanistan, Albania, Algeria, Azerbaijan, Bahrein, Bangladesh, Brunei, Chad, Djibouti, Egypt, Gambia, Guinea, Indonesia, Iran, Iraq, Jordan, Kazakhstan, Kuwait, Kyrgyzstan, Lebanon, Libya, Malaysia, Maldives, Mali, Mauritania, Mayotte, Morocco, Niger, Oman, Pakistan, Qatar, Saudi-Arabia, Senegal, Somalia, Sudan, Syria, Tajikistan, Tunisia, Turkey, Turkmenistan, Uzbekistan, Yemen, Sierra Leone, United Arab Emirates, Western Sahara, and Burkina Faso.

information on the number of all charges that the suspect is facing in a given criminal case within which the observed charge is processed. The average number of charges equals 2.4. Finally, we know the crime classification of the most severe among the charges faced by the suspect in the given case (in the Netherlands, sentencing of convicted offenders is by law based on the most severe charge). Violent crimes in our primary sample are subdivided into the categories physical abuse, physical threat, assault, crime against life, use of weapon, and other violent crime. Physical abuse, threat, and crime against life together account for 82 percent of the charges included in our primary sample.

4. Methods and Baseline Results

4.1. Empirical Strategy

To investigate the impact of van Gogh's murder on prosecution of suspects born in Muslim-majority countries, we exploit the fact that the van Gogh murder was plausibly exogenous to prosecutorial decision-making and use a difference-in-differences approach. In our primary approach we focus exclusively on the set of violent criminal cases that were processed by the Amsterdam-district prosecutor office during the 12-week interval surrounding van Gogh's death (i.e. six weeks prior to and six weeks after the murder) and involve a male suspect. We center our attention on violent crimes entailing a male suspect because we anticipate those crimes to especially likely invoke a connotation in the prosecutors' minds with the van Gogh murder. We restrict the geographic dimension of our primary sample to the cases processed by the Amsterdam prosecutors' office because it is the Amsterdam region that was most immediately impacted by van Gogh's killing. To the extent that the murder impacted prosecutorial decision-making, we expect the effect to be detected most clearly in the Amsterdam region.

Finally, the choice of the 12-week time interval implies a one-and-a-half-month time period after the van Gogh murder during which the event was debated most actively in the media and the society. To the extent that van Gogh's murder and the subsequent reaction of the Dutch public shaped prosecutors' views, a six-week post-murder period implies a reasonably wide time interval that, on the one hand, facilitates the use of a sufficiently large estimating sample and, at the same time, allows for any impact on prosecutorial decision-making to be detectable in the data. In the section on robustness checks below, we examine the sensitivity of our results to all of the above-noted choices.

Our empirical approach presumes that the suspect and crime characteristics in the sample of investigated cases do not vary systematically with the timing of when the corresponding cases were processed relative to November 2, 2004, the day of the van Gogh murder. To assess the validity of this assumption, we estimate a series of ordinary-least-squares (OLS) regressions, one per suspect or other case characteristic. In each regression, the explanatory variable is the indicator equal to one if the pertinent case was processed after the van Gogh murder and zero otherwise. Intuitively, whether the case was processed prior to or after the van Gogh murder should not be a statistically significant predictor of suspect and other case characteristics. Table 3 shows this is indeed the case. Cases processed prior to the van Gogh murder are not systematically different from cases processed after the murder. In Section 5 we further address the issue of case composition by presenting estimates using only the sample of crimes committed before the van Gogh killing.

We thus posit the following linear probability model:

$$\Pr(\textit{Prosecute}_{crst} = 1 | x_{rst}) = \alpha + \beta \textit{Post}_t \times \textit{Muslim}_s + \gamma \textit{Muslim}_s + \lambda_t + \delta' X_{rst} + \kappa' \textit{Post}_t \times X_{rst}, \quad (1)$$

where the unit of observation is a charge c connected to criminal case r that involves suspect s and is processed by the prosecution office during time period t . In expression (1), $\textit{Prosecute}_{crst}$ is a binary indicator equal to one if the applicable prosecution office chooses to prosecute the pertinent charge and zero otherwise. x_{rst} is the vector of observables included on the right-hand side of (1). \textit{Post}_t is a binary indicator equal to one if the prosecution office processed the charge during a time period after the van Gogh murder and zero if the charge was processed before the van Gogh murder. \textit{Muslim}_s is a binary indicator equal to one if suspect s , who has been charged with charge c , had been born in a Muslim-majority country and zero otherwise. λ_t is the fixed effect for a specific biweek (i.e. 14-day period) before or after the van Gogh murder. Given our focus on a 12-week bandwidth around the time of the murder, we therefore include five different biweek dummies. We omit the dummy for the biweek immediately preceding the van Gogh murder. The biweek fixed effect λ_t fully absorbs the timing of the processing of the applicable charge relative to the timing of the van Gogh murder, rendering a separate inclusion of \textit{Post}_t on the right-hand side of (1) redundant.¹⁰

¹⁰ In reporting the baseline results below, we also show the results based on the simplest difference-in-differences specification in which the biweek fixed effects are replaced with the Post Murder (\textit{Post}_t) dummy.

X_{rst} is the vector of suspect and crime level controls. Among suspect-level controls, we include the full set of indicators for the suspect's age group and labor market status at the time of the prosecution decision. We further control for whether the suspect has a prior criminal history. Among the crime level controls, we control for whether the suspect faced multiple charges; the fixed effect for the type of the most severe charge, which defines the crime type associated with the particular charge under consideration; and the crime month-by-year fixed effects (not shown in Tables 1 and 2), where month and year pertain to the timing of the committed crime (as opposed to the timing of processing of the criminal case by the prosecution office). The latter set of fixed effects is intended to address any impact on prosecutorial behavior of seasonality in crime.¹¹

Lastly, $Post_t \times X_{rst}$ is the full set of interactions between the $Post_t$ indicator for the timing of the processing of the suspect's charge relative to the timing of van Gogh murder and the subset of suspect and crime level controls. The key reason for the inclusion of the full set of the resulting interaction effects is to ensure that the any estimate of the coefficient on the interaction term between $Post_t$ and $Muslim_s$ is not spuriously reflecting the impact of some suspect or case level factor other than the suspect hailing from a Muslim-majority country.¹² At the same time, the van Gogh murder may have affected prosecutorial processing of charges along characteristics that are independent of whether the suspect originates from a Muslim-majority country. If so, the inclusion of the $Post_t \times X_{rst}$ interactions may further reduce the residual variance and thereby sharpen the estimate of the key coefficient of interest.

The focal coefficient of interest in expression (1) is β , the difference-in-differences (DD) estimate of the impact of the van Gogh murder on the prosecutorial treatment for unrelated violent crimes of suspects born in a Muslim-majority country. More precisely, β captures the difference between (i) the change in the probability of prosecution following the van Gogh murder for charges involving male suspects born in Muslim-majority countries and (ii) the analogous change for charges involving other male suspects. That is, the difference (i) contrasts the pre- versus post-murder prosecutorial decisions for male suspects born in Muslim-majority countries. However, we do not observe the counterfactual of how these suspects would have been treated by the prosecutors

¹¹ See Philippe (2017) for an analogous approach in the context of judicial decision-making. There are on average 33 observations per crime month-by-year cell in our data. The mean standard deviation of the focal $Post_t \times Muslim_s$ interaction variable across the crime month-by-year cells is 0.141. Thus, there is ample variation in our data to facilitate the estimation of a model with crime month-by-year fixed effects.

¹² See, for example, Jia and Nie (2017) for an analogous empirical strategy in the implementation of the difference-in-differences approach in a non-legal context.

shortly after November 2, 2004 had the van Gogh murder not occurred. To construct the missing counterfactual, we therefore use the change in the probability of prosecution for charges involving other male suspects who are either Dutch-born or, if they are foreign-born, were not born in a Muslim-majority country. The latter set of suspects constitutes our control group. (In Section 5.5 we demonstrate the robustness of our findings to varying the definition of treated and control group.) Subtracting difference (ii) from difference (i) then provides a DD estimate of the effect of the van Gogh assassination on the likelihood of prosecution of male suspects of violent crimes who were born in a Muslim-majority country.

In order for suspects who were not born in a Muslim-majority country to serve as a good control group for suspects who were born in a Muslim-majority country, in the absence of the van Gogh murder and conditional on controls the change in the prosecution prospects would have to be the same for both groups. (We discuss placebo treatments in Section 5.4.) Figure 2 indicates this is indeed the case, implying that a key identifying assumption underpinning our DD approach is plausibly satisfied in our data. (In Section 5, we provide further evidence in favor of this point by demonstrating the absence of pre-trends in prosecutorial decision-making based on estimation of multiple dynamic specifications.) In the six weeks prior to the van Gogh murder, the unconditional average likelihood of prosecution by the Amsterdam prosecution office was decreasing for both suspects born in a Muslim-majority country and other suspects. Suspects from Muslim-majority countries were in general prosecuted less likely than other suspects, a pattern consistent with the perception of those who had argued that the Amsterdam public administration had a pro-minority stance (see Section 2.1).

Figure 2 further suggests that the van Gogh murder induced a sharp, but temporary, deviation from this trend for suspects born in Muslim-majority countries. After the murder, the average unconditional probability of prosecution for charges involving this group of suspects increases for about a month, and then falls to a level similar to that before the killing. In contrast, the average unconditional probability of charge prosecution for other suspects during the first four weeks following the murder remains at lower or roughly similar levels as prior the murder. The average prosecution probability increases somewhat in the fifth and sixth week after the murder, but not above the level observed in five to six weeks prior to the murder.

Our data do not allow us to observe the identity of the prosecutor that handled a particular charge. We thus do not control for prosecutor fixed effects. However, this should come without

loss of generality. The institutional features of the Dutch prosecution process and prosecutorial case assignment (see Section 2.2) have not been subject to change during the time period of our focus. Thus, conditional on the crime type, the unobservable prosecutor characteristics should be orthogonal to the suspect's ethnoreligious identity and the timing of case processing. In Section 5 we provide evidence in support of this point by assessing the sensitivity of our estimates to omitted variable bias.

Our use of a linear probability model by default gives rise to heteroscedastic errors. In addition, it is plausible that the unobservables affecting the decision of a prosecution office concerning whether to prosecute a charge are correlated, which suggests the need for clustering of standard errors. The choice of the level at which clustering should be applied, however, is not obvious. Because we do not observe the identity of the prosecutors, we are unable to cluster at the prosecutor level. Irrespective of that, we would expect the unobservables to be correlated primarily across charges (if multiple) stemming from a given single-suspect criminal case; the assignment to a treatment is in our context most naturally thought of as clustered at the level of a case (Abadie et al. 2017). As our primary approach we thus base statistical inference on heteroscedasticity-robust standard errors that are clustered at the level of a (single-suspect) criminal case, an approach resulting in 766 clusters. In addition, for our baseline estimates we also show the results when we cluster standard errors at the level of the city of crime (giving rise to 23 clusters) and the accused's country of birth (amounting to 49 clusters), respectively. As we clarify below, none of these alternative clustering approaches in any way alter our results; in fact, our primary approach of clustering at the level of the case turns out to be the most conservative among the considered options. Finally, we have verified that none of our findings change if we cluster standard errors at the level of the suspect or the week of processing of the crime.¹³ In sum, our findings are not sensitive to alternative clustering approaches.

4.2. Baseline Results

The results based on our main specification are reported in Table 4. In addition to our primarily clustering approach (at case level) we show the estimated standard errors and indicate statistical significance based on clustering at the city of crime and the accused's country of birth, respectively. Column (1) shows the results for the basic difference-in-differences specification where we replace

¹³ Detailed results corroborating this point are available upon request.

the biweek fixed effects with the Post Murder ($Post_t$) dummy. Based on the estimates in column (1), relative to the situation before the murder, after the van Gogh murder male suspects from Muslim-majority countries are on average charged 16.4 percentage points more likely. The effect is statistically significant irrespective of the clustering approach.

In column (2), we replace the Post Murder dummy with biweek fixed effects. In columns (3)-(7) we then progressively include a more and more encompassing set of fixed effects and controls. In column (3), we add other suspect characteristics. In column (4), we augment the specification in column (3) by controlling for other criminal case characteristics. In column (5), we include the crime month-by-year fixed effects. Finally, in column (6) we add a full set of interactions between the indicator for the post-murder period and all of the suspect and case controls. Based on the resulting preferred specification, as a consequence of the van Gogh murder, the prospects of the Amsterdam prosecution office pursuing charges against male suspects from Muslim-majority countries increased by 19.7 percentage points. The resulting effect is statistically significant at five percent significance level using our primary clustering approach and at one percent significance level using alternative clustering approaches shown in Table 4. The effect is also quantitatively noteworthy (before the van Gogh murder, the mean unconditional likelihood of charging a suspect born in a Muslim-majority country was equal to 0.6827). To put our estimate in a further context, Brodeur and Wright (2019), for example, find that the September 11, 2001 terrorist attacks resulted in a 3.2 percentage point decrease in the prospects of U.S. asylum judges granting asylum to applicants from Muslim-majority countries (the reported unconditional probability of an applicant being granted asylum is 0.199). All in all, our estimated effect of van Gogh's murder on the likelihood of prosecution of suspects born in Muslim-majority countries is therefore considerable in size.

Finally, the estimate of the coefficient on the Muslim dummy is consistently negative and statistically significant based on our preferred specification. Congruent with the pattern shown in Figure 2, prior to the van Gogh killing, charges for violent crimes faced by male suspects from Muslim-majority countries were in Amsterdam prosecuted on average somewhat less likely than equivalent charges faced by other suspects.

5. Sensitivity Analysis and Further Results

We subject our results to a series of robustness checks and alternative model specifications. We first investigate the robustness of our findings to adding group-specific time trends, ascertain

dynamic effects, and further assess the presence of any undesirable pre-trends in the data. We then explore the sensitivity of our results to a range of exercises. We vary the time bandwidth. We alter the definition of the treated group. We perform multiple falsification tests. We restrict the sample only to crimes that were committed before the van Gogh murder. We quantify the sensitivity of our baseline estimates to omitted variable bias. We explore the results for crimes other than violent crimes and for prosecution offices other than the Amsterdam office. Finally, we examine the results using an alternative outcome variable. Throughout the remaining analysis, we base inference on standard errors clustered at the case level, our primary and most conservative approach.

5.1. Adding Suspect Group-Specific Time Trend

In addition to the visual assessment of the parallel trend assumption (Figure 2), an alternative check for the appropriateness of the DD identification strategy entails the estimation of a model featuring time trends specific to the treated and the control groups, respectively. We thus estimate a specification that includes among the set of covariates in model (1) a linear time trend for treated-group suspects (born in Muslim-majority country) versus control-group suspects, thereby allowing the two suspect groups to follow different trends with regard to prosecutorial decisions. Column (7) of Table 4 shows the results based on the resulting specification. The DD estimate of the coefficient of interest continues to be positive and statistically significant, while the coefficient on the treated group-specific time trend is statistically insignificant (p -value equals 0.262). This is further evidence in support of our identification strategy.

5.2. Ascertaining Dynamic Effects and Assessing Pre-Trends

Our baseline estimates evaluate the average effect of the van Gogh murder. We next allow for a more flexible specification and evaluate the dynamic effects. The resulting approach also allows us to test for any unwanted pre-trends in the data. To this end, we posit the following dynamic version of model (1):

$$\Pr(\text{Prosecute}_{crist} = 1 | x_{rst}) = \alpha + \sum_{\tau} \beta_{\tau} \text{Biweek}_{\tau} \times \text{Muslim}_s + \gamma \text{Muslim}_s + \lambda_t + \delta' X_{rst} + \kappa' \text{Post}_t \times X_{rst}, \quad (2)$$

where the terms $\text{Biweek}_{\tau} \times \text{Muslim}_s$ are the interactions between the indicator for a Muslim suspect and different biweek dummies, with $\tau \in \{-3, -2, 1, 2, 3\}$. The omitted biweek is thus the biweek just

before the van Gogh murder ($\tau=-1$). The remaining elements of expression (2) are as defined under expression (1).

The results based on the estimation of model (2) are presented in Table 5. Column (1) shows the results with core suspect and crime controls. In column (2) we add crime month-by-year fixed effects. In column (3), the most comprehensive and our preferred specification, we further include the interactions between the Post Murder dummy and the controls. The estimates show that the van Gogh murder-induced positive effect on the likelihood of charging suspects born in Muslim-majority countries is limited to the first two biweeks (i.e. a month) after the murder, and dissipates thereafter. The point estimate of the effect in the first biweek following the murder exceeds that during the second biweek after the murder. The difference between the two effects is marginally statistically significant (p -value for the test of equality of pertinent coefficients equals 0.060). The effect during the third biweek after the van Gogh murder is statistically insignificantly different from zero. In addition, the estimates in Table 5 show that our data exhibit no pre-trends in prosecutorial decision-making in the time period prior to the van Gogh murder. The dynamic effects, implied by the estimates in Table 5, are illustrated graphically with the aid of Figure 3.

5.3. *Varying the Bandwidth*

We next assess the sensitivity of our results to varying the time period of the cases included in the estimated sample, i.e. the bandwidth. While our baseline estimates use a 12-week bandwidth, we examine the estimates based on both narrower and wider bandwidths. The temporal dynamics of the van Gogh murder effect highlighted in the previous subsection suggest that the effect is concentrated in the first month after the murder. Thus, we would expect the narrowing of the bandwidth to a time period shorter than 12 weeks to increase the magnitude of the implied average effect. On the other hand, the widening of the bandwidth to a time period longer than 12 weeks should lead to a dilution of the average effect, until the effect can no longer be detected in the data.

We first vary the bandwidth using model specification (1). We investigate 8-week, 10-week, 14-week, 16-week, and 20-week bandwidth. The corresponding results, reported in Table 6, are consistent with our supposition. Using bandwidths that are narrower than the 12-week bandwidth used to generate our baseline estimates increases somewhat the estimated average effect of the van Gogh murder. On the other hand, using bandwidths that are wider than the baseline 12-week bandwidth reduces both the magnitude and the statistical significance of the estimated effect. With a 16-week and 20-week bandwidth, the effect of the van Gogh murder is no longer detectable.

We also vary the bandwidth using dynamic model specification (2), focusing on 8-week, 16-week, and 20-week bandwidths. The results are reported in Table 7. The effect continues to be concentrated in the first month after the murder. Notably, we continue to see no pre-trends in prosecutorial decision-making.

5.4. Placebo (Falsification) Tests

We perform a series of placebo tests. Drawing on the sample used to generate the results in Table 4, we first re-assign the suspect characteristic of having been born in a Muslim-majority country, attributed to suspects in 17.7 percent of the charges in our primary sample, to randomly chosen charges and re-estimate the version of model (1) as featured in column (6) of Table 4. We perform 500 independent iterations of such re-assignment and model re-estimation. Figure 4 summarizes the corresponding estimates of the focal coefficient β and the corresponding t -statistics. Intuitively, if our estimates reported in Table 4 truly capture the effect of the van Gogh murder based on suspects' birth in a Muslim-majority country as opposed to some other unobserved suspect characteristics, then we would expect the distribution of the resulting estimates and t -statistics to be centered at zero. Figure 4 confirms that this is indeed the case.

We also conduct placebo tests for cases processed within six weeks before or after November 2 of years 2001, 2002, and 2003, respectively. Because the van Gogh murder took place on November 2 of 2004, the DD estimate should not be statistically significant from zero for samples based on earlier years. The results shown in Table 8 confirm this is indeed the case. That is, our estimates discussed in Section 4 are not an artifact of ethnoreligious-time effects in prosecution that are unrelated to the van Gogh murder per se but somehow naturally tend to occur on November 2.

5.5. Varying Definition of Treated and Control Groups

We re-estimate the baseline specification featured in column (6) of Table 4 while varying the definition that a suspect was born in a Muslim-majority country. Specifically, we gradually increase the threshold value of the proportion of population in the country of origin of the suspect that self-identifies as Muslim from the baseline 50 percent to 90 percent in ten-percentage-point increments. The results are reported in columns (1) through (4) of Table 9. The DD estimate of the effect of van Gogh murder on the prospects of charging suspects born in Muslim-dominated countries continues to be positive and statistically significant.

We re-define the notion of our treated group to include not only suspects who were themselves born in a Muslim-majority country but also suspects whose parents had been born in a Muslim-majority country. The corresponding treated group, which now encompasses 311 charges (26.1 percent of our primary sample), therefore includes suspects who were themselves born in the Netherlands but whose parents had been born in a country where at least 50 percent of the population self-identifies as Muslim. The results based on the corresponding model specification are reported in column (5) of Table 9. The estimate of the coefficient of interest is positive and statistically significant.

We explored a specification where the treatment group indicator takes on the value one only if the suspect's country of birth is Morocco, the country of origin of the family of van Gogh's assassin. The results are shown in column (6) of Table 9. The estimated effect of interest is not statistically significant from zero. We thus do not find evidence that the van Gogh murder affected the prospects of charging Moroccan suspects specifically. That is, the effect of the van Gogh murder on prosecutorial decision-making that we have uncovered pertains to suspects who were themselves born in a Muslim-majority country or whose parents had been born in Muslim-majority country.

We re-define our control group to consist only of suspects who were not born abroad and whose parents were also not born abroad. We do so because the extraneous event of our focus perhaps affected not only the prosecution of suspects born in Muslim-majority countries, but also exhibited spillover effects on the prosecution of other foreign-born suspects (i.e. those not born in a Muslim-majority country). In contrast, the prosecution of 'indigenous' Dutch suspects would have been least likely in any way impacted by the van Gogh killing. The parallel trends figure (not reported) for the resulting sample looks qualitatively very similar to Figure 2. The estimation results are reported in column (7) of Table 9. The DD estimate of interest is still positive and statistically significant.

5.6. Restricting the Sample to Crimes Committed Before the van Gogh Murder

The covariate balancing tests (Section 4) demonstrate that the charges processed prior to the van Gogh murder are not systematically different from the charges processed after the murder. However, a concern may be that the crimes committed before the van Gogh murder might still differ systematically from the crimes committed after the van Gogh murder along some unobservable dimension not adequately captured by our controls, and thus that our results reflect

such unobserved differences in the nature of committed crimes, as opposed to differences in prosecutorial decision-making per se. To address this concern, we re-estimated our preferred model specification (1) using only the sample of charges pertaining to crimes committed before the van Gogh murder, but processed by the prosecutorial office either before or after the van Gogh murder.

Restricting the estimating sample to crimes committed prior to van Gogh murder also addresses the concern that, after the van Gogh murder, violent crimes entailing male Muslim suspects were somehow subject to a different law enforcement or prosecutorial regime. For example, the reader might worry that, after the van Gogh killing, the police deliberately decided to target Muslim minorities more aggressively or that, despite well-defined institutional rules for assignment of cases to prosecutors (see Section 2.1), cases involving Muslim suspects nevertheless somehow ended up being allocated to the harshest prosecutors. If true, then our baseline results could be merely a consequence of altered law enforcement focus or changes in prosecutorial case-assignment regime as opposed to changes in prosecutorial treatment of Muslim suspects per se. Focusing on the sample of crimes committed before van Gogh's assassination by definition precludes the potentially confounding effect of altered police behavior. In addition, because cases for crimes committed prior to the van Gogh murder were surely allocated to prosecutors before the van Gogh murder, restricting the sample to crimes committed before van Gogh's assassination mitigates the concern that our findings could be driven by an ad-hoc, formally-never-approved change in prosecutorial case assignment.

Table 10 thus reports the results for specifications analogous to those in columns (2) through (6) of Table 4, as well as for the dynamic specification reported in Table 5, but only for the sample of crimes committed before the van Gogh murder. The results in Table 10 closely resonate with the results reported in Tables 3 and 4. That is, even when restricting the sample to only the set of crimes committed before the van Gogh murder, we continue to see that the van Gogh murder increased the prospects of prosecution of Muslim-minority suspects.

5.7. Quantifying Sensitivity to Omitted Variable Bias

Evidence of covariate balance before versus after the van Gogh murder (Table 3) and the estimation results based on restricting the sample exclusively to crimes committed prior to the van Gogh murder (Section 5.6) alleviate the concern that our results are perhaps driven by unobservables. As a further means of addressing the concern that our findings could be explained

by omitted factors that affect prosecutorial decisions, vary across Muslim and non-Muslim suspects, and change at the time of the van Gogh killing, we use the method developed by Oster (2019) and apply it to our preferred baseline specification in column (6) of Table 4.

Oster formalizes the intuitive notion that, in a linear regression, knowledge about the importance of observed controls can be used to assess the bias from unobserved controls. Oster's approach involves making an assumption about the magnitude of R_{max} , the value of R^2 obtained on the basis of a hypothetical regression containing all relevant (observed and unobserved) controls. Oster's suggested value of R_{max} equals $\min\{1.3R^2, 1\}$. We follow her recommendation but also explore the robustness of our findings to using even more conservative approaches, thus additionally setting R_{max} to equal $\min\{2R^2, 1\}$ and $\min\{4R^2, 1\}$, respectively.

Oster's method can be used to evaluate the value of δ , the ratio of the importance of selection on unobservables versus observables that would render the estimate of interest equal to zero. Part A of Table 11 summarizes our results based on the corresponding exercise. Under Oster-suggested value of R_{max} (column (1)), selection on unobservables would have to be more than 4.3 times as important as selection on observables to explain away the effect of interest. It is only under the most conservative among the considered scenarios (column (3)) that selection on unobservables would not need to be at least as important than selection on observables to explain away the effect of interest.

Oster's method can alternatively be used to estimate the bounds on the estimated effects for a given magnitude of δ . Once more following Oster's recommendation, we set $\delta=1$ (i.e. equal importance of selection on unobservables versus observables). Part B of Table 11 summarizes the results. Regardless of the assumed value of R_{max} , the estimated coefficient bounds never include the value zero, an indication that, under the stated assumptions, omitted variables cannot explain the estimated DD effect. In sum, the application of Oster's (2019) methodology shows that our estimates are highly unlikely an artifact of omitted variable bias.

5.8. Probing the Mechanisms: Results by Crime Type and for Other Prosecutor Offices

Our baseline estimates focus on violent crimes that involve male suspects and that were processed by the Amsterdam prosecution office. To the extent that the van Gogh murder indeed shaped prosecutorial decision-making, the corresponding set of crimes is in theory most likely to have resulted in a differential treatment of Muslim suspects. But is the effect of the van Gogh murder, as evidenced by the results reported in previous sections, also detectable with respect to other types

of crimes committed in Amsterdam or with respects to an analogous set of crimes that fall under the jurisdiction of other prosecution offices in the Netherlands? Answers to these questions can help us distinguish between emotions-based versus other mechanisms that can explain our results (see Section 2.3).

We first examine the role of crime type. To this end, we re-estimate model (1) using the sample of cases processed by the Amsterdam prosecution office and entailing a male suspect, but this time involving non-violent crimes. Specifically, we classify the corresponding crimes under four headings: property crimes (encompassing offences such as theft, burglary, scams, and money laundering), drug-related crimes (possession and distribution), destruction and violation of public order (due to destruction of private and public property, trespassing, non-compliance with an official order), and traffic offences (e.g. driving under influence, hit and run, driving under a driving ban or with an illegal license plate). We then re-estimate the specification as featured in column (6) of Table 4 separately for each set of crimes. Table 12 displays the results. We find no effect of the van Gogh murder on the prospects of charging Muslim suspects for the respective categories of crime designated as property, drugs, and destruction and violation of public order (columns (1) through (3)).

Interestingly, we find that the van Gogh murder actually decreased the likelihood of charging Muslim suspects in traffic offences (column (4)). One possible interpretation of this result may be that, in the presence of scarce prosecutorial resources, the van Gogh crime induced a substitution of prosecutorial effort. Based on this argument, following the van Gogh murder, the Amsterdam prosecutors (intentionally or unintentionally) zeroed in on Muslim-minority suspects accused of violent crimes. The side-effect of this, however, may have been laxer prosecution of Muslim suspects in the context of lesser offences, such as those committed with regard traffic issues.

We next examine whether our findings for the Amsterdam prosecution office can also be detected in the context of other prosecution offices in the Netherlands. We thus re-estimate model (1) using cases processed by other prosecution offices in the Netherlands. To render the resulting analysis comparable with the baseline estimates shown in column (6) of Table 4, we continue to focus on violent crimes entailing a male suspect. Table 13 shows the results, with each column pertaining to a different prosecutors' office. We find no evidence of the van Gogh murder effect for any of the nine non-Amsterdam public prosecution offices. Thus, the documented impact of

the van Gogh murder on prosecution of Muslim-minority suspects is detectable only in the context of the prosecutorial processing of crimes that occurred in the geographic region of the murder.

Overall, the results in Tables 12 and 13 appear to be inconsistent with a purely emotions-based explanation of our results. If the van Gogh assassination had spurred especially strong feelings of anger and animosity toward Muslim-minority men, we should have observed more aggressive prosecution of the corresponding suspects across multiple crime types and across other Dutch prosecution offices. This is not what our data show. But our estimation results per se do not allow us to distinguish between the availability heuristic-based and signaling explanations. On this point, we only offer our conjectural thoughts. As noted, there is no indication that, following the murder, Amsterdam prosecutors faced any kind of top-down pressure to prosecute specific suspect groups more aggressively, or that the prosecutors perhaps strove to display their toughness directly to the Muslim community (see Section 2.3). There is of course the possibility that such motives existed, but were kept private. If so, then given the extraordinarily high profile of the event, after 17 years, some of the underlying narratives would have presumably surfaced in the public. We have encountered no such narratives. Thus, on institutional grounds, we view the signaling story as a possible, but likely not the predominant explanation for our findings. Availability heuristic-induced unconscious perceptions, followed by prosecutors' conscious choice to rely on the corresponding mental shortcut, would seem at least as persuasive of an explanation for our results.

5.9. Using an Alternative Outcome Variable

We re-estimated our main model specification upon replacing the binary indicator $Prosecure_{crst}$ with a binary indicator $Drop_{crst}$ equal to one if charge c in criminal case r against suspect s processed during time period t was altogether dropped, and zero otherwise. The results based on a set of specifications analogous to that reported in columns (2) through (6) of Table 4 are shown in Table 14. According to the specification with the full set of controls in column (5), after the van Gogh murder, charges against suspects born in Muslim-majority countries are on average dropped 6.9 percentage points less likely. (Before the murder, the mean unconditional probability that charges against a suspect born in a Muslim-majority country would be dropped equals 0.0673.) The effect is statistically significant. Congruent with the results based on the main specification, these estimates imply that the van Gogh murder led to a harsher prosecutorial treatment of Muslim-minority suspects.

6. Conclusion

We have drawn on the universe of criminal cases processed by the Dutch public prosecution offices shortly before and after Mohammed Bouyeri, a Moroccan-Dutchman who had become an Islamic extremist, in Amsterdam brutally murdered Theo van Gogh, a Dutch filmmaker notorious for his critique of Islam. The data reveal that the murder exerted an effect on the prosecution for entirely unrelated crimes of male suspects born in Muslim-majority countries. Specifically, after van Gogh's assassination, the prospects of prosecution of charges linked to violent crimes committed in the Amsterdam area and involving male suspects born in Muslim-majority countries on average increased by more than 19 percentage points. Based on a dynamic specification, the effect is discernible only during the first month following the van Gogh murder and disappears thereafter. We do not find an effect of the murder for non-violent crimes and for violent crimes processed by other public prosecution offices outside of Amsterdam.

These empirical findings are in principle consistent with both a signaling story, whereby career-motivated prosecutors signal their toughness to their political-administrative principals or the male Muslim community at large, and a behavioral explanation, under which the prosecutors' murder-induced unconscious perceptions of culpability of male Muslim suspects for violent crimes (an availability heuristic) were followed by the prosecutors' conscious response to rely on the corresponding mental shortcut. On institutional grounds, the latter explanation appears somewhat more plausible than the former explanation. But our results do not support the hypothesis that the observed effects are a consequence of the prosecutors' negative emotional response.

Our analysis hence shows that salient extraneous events can fundamentally shape not only the decisions of judges and juries, as demonstrated by prior literature, but also of prosecutors as crucial but comparatively understudied decision-makers in the administration of criminal justice. Prosecutors select the subset of the criminal cases to be tried in court, yet studies of judicial and jury behavior are rarely able to explicitly address the corresponding selection issues and unambiguously pinpoint any extralegal disparities to the decisions of specific actors. In demonstrating that prosecutorial decisions can be an important source of bias in the processing of criminal cases, our results are thus of relevance also to scholars of judicial and jury behavior. More generally, our findings shed novel light on the ongoing policy and academic debates about the importance of extralegal factors for justice system outcomes.

References

- Abadie, Alberto, Susan Athey, Guido W. Imbens, and Jeffrey M. Wooldridge. 2017. "When Should You Adjust Standard Errors for Clustering?" NBER Working Paper No. 24003.
- Bandyopadhyay, Siddhartha and Bryan McCannon. 2014. "The Effect of the Election of Prosecutors on Criminal Trials." *Public Choice*, 161:1, 141-156.
- Bjerk, David, 2005. "Making the Crime Fit the Penalty: The Role of Prosecutorial Discretion under Mandatory Minimum Sentencing." *Journal of Law and Economics*, 48:2, 591-625.
- Boylan, Richard T. 2004. "Salaries, Turnover, and Performance in the Federal Criminal Justice System." *Journal of Law and Economics*, 47:1, 75-92.
- Boylan, Richard T. and Cheryl X. Long. 2005. "Salaries, Plea Rates, and the Career Objectives of Federal Prosecutors." *Journal of Law and Economics*, 48:2, 627-651.
- Brodeur, Abel and Taylor Wright. 2019. "Terrorism, Immigration and Asylum Approval." *Journal of Economic Behavior and Organization*, 168:C, 119-131.
- Burns, Kylie. 2016. "Judges, 'Common Sense' and Judicial Cognition." *Griffith Law Review*, 25:3, 319-351.
- Cohen, Alma and Crystal S. Yang. 2019. "Judicial Politics and Sentencing Decisions." *American Economic Journal: Economic Policy*, 11:1, 160-191.
- Danziger, Shai, Jonathan Levav, and Liora Avnaim-Pesso. 2011. "Extraneous Factors in Judicial Decisions." *Proceedings of the National Academy of Sciences*, 108:17, 6889-6892.
- De Haan, Willem. 2007. "Morele Paniek of Cultureel trauma? Over de Betekenis en de Gevolgen van de Moord op Theo van Gogh." *Tijdschrift voor Criminologie*, 50:3, 252-264.
- Depew, Briggs, Ozkan Eren, and Naci Mocan. 2017. "Judges, Juveniles, and In-Group Bias." *Journal of Law and Economics*, 60:2, 209-239.
- Diamond, Shari Seidman and Loretta J. Stalans. 1989. "The Myth of Judicial Leniency in Sentencing." *Behavioral Sciences and The Law*, 7:1, 73-89.
- Dyke, Andrew. 2007. "Electoral Cycles in the Administration of Criminal Justice." *Public Choice*, 133:3/4, 417-437.
- Eren, Ozkan, and Naci Mocan. 2018. "Emotional Judges and Unlucky Juveniles." *American Economic Journal: Applied Economics*, 10:3, 171-205.
- Fischman, Joshua B. and Max M. Schanzenbach. 2012. "Racial Disparities Under the Federal Sentencing Guidelines: The Role of Judicial Discretion and Mandatory Minimums." *Journal of Empirical Legal Studies*, 9:4, 729-764.
- Finseraas, Henning, Niklas Jakobsson, and Andreas Kotsadam. 2011. "Did the Murder of Theo van Gogh Change Europeans' Immigration Policy Preferences?" *Kyklos*, 64:3, 396-409.
- Garoupa, Nuno. 2012. "The Economics of Prosecutors." In: Harel, Alon and Keith N. Hylton (eds.), *Research Handbook on the Economics of Criminal Law*, Northampton, MA: Edward Elgar Publishing, 231-242.
- Gautier, Pieter A., Arjen Siegmans, and Aico Van Vuuren. 2009. "Terrorism and Attitudes Towards Minorities: The Effect of the Theo van Gogh Murder on House Prices in Amsterdam." *Journal of Urban Economics*, 65:2, 113-126.

- Gilliéron, Gwladys. 2014. *Public Prosecutors in the United States and Europe. A Comparative Analysis with Special Focus on Switzerland, France, and Germany*. New York, NY: Springer International Publishing.
- Glaeser, Edward L., Daniel P. Kessler, and Anne Morrison Piehl. 2000. "What Do Prosecutors Maximize? An Analysis of the Federalization of Drug Crimes." *American Law and Economics Review*, 2:2, 259-290.
- Hajer, Maarten and Justus Uitermark. 2008. "Performing Authority: Discursive Politics after the Assassination of Theo van Gogh." *Public Administration*, 86:1, 5-19.
- Heyes, Anthony and Soodeh Saberian. 2019. "Temperature and Decisions: Evidence from 207,000 Court Cases." *American Economic Journal: Applied Economics*, 11:2, 238-265.
- Jia, Ruixue and Huihua Nie. 2017. "Decentralization, Collusion, and Coal Mine Deaths." *Review of Economics and Statistics*, 99:1, 105-118.
- Korobkin, Russell B. and Thomas S. Ulen. 2000. "Law and Behavioral Science: Removing the Rationality Assumption from Law and Economics." *California Law Review*, 88:4, 1051-1144.
- Kuran, Timur. 1995. *Private Truths, Public Lies. The Social Consequences of Preference Falsification*. Cambridge, MA: Harvard University Press.
- Lerner, Jennifer S., Ye Li, Piercarlo Valdesolo, and Karim S. Kassam. 2015. "Emotion and Decision Making." *Annual Review of Psychology*, 66, 799-823.
- Lim, Claire S.H., Bernardo S. Silveira, and James M. Snyder. 2016. "Do Judges' Characteristics Matter? Ethnicity, Gender, and Partisanship in Texas State Trial Courts." *American Law and Economics Review*, 18:2, 302-357.
- Luna, Erik and Marianne Wade (eds.) 2012. *The Prosecutor in Transnational Perspective*. Oxford, UK: Oxford University Press.
- McCannon, Bryan C. 2013. "Prosecutor Elections, Mistakes, and Appeals." *Journal of Empirical Legal Studies*, 10:4, 696-714.
- McCannon, Bryan C. and Mark Wilson. 2019. "'A Million Dollars in Free Advertising' Politics and Sex Offense Prosecution in the Wake of Duke Lacrosse." *Supreme Court Economic Review*, 27, 167-187.
- Nadel, Melissa R., Samuel J. A. Scaggs, and William D. Bales. 2017. "Politics in Punishment: The Effect of the State Attorney Election Cycle on Conviction and Sentencing Outcomes in Florida." *American Journal of Criminal Justice*, 42:2, 845-862.
- Oster, Emily. 2019. "Unobservable Selection and Coefficient Stability: Theory and Evidence." *Journal of Business and Economic Statistics*, 37:2, 187-204.
- Philippe, Arnaud. 2017. "Does Introducing Lay People in Criminal Courts Affect Judicial Decisions? Evidence from French Reform." *International Review of Law and Economics*, 52, 1-15.
- Philippe, Arnaud and Aurélie Ouss. 2018. "'No Hatred or Malice, Fear or Affection': Media and Sentencing." *Journal of Political Economy*, 126:5, 2134-2178.
- Rasul, Imran and Brendon McConnell. 2020. "Contagious Animosity in the Field: Evidence from the Federal Criminal Justice System." *Journal of Labor Economics*, forthcoming.
- Rehavi, M. Marit and Sonja B. Starr. 2014. "Racial Disparity in Federal Criminal Sentences." *Journal of Political Economy*, 122:6, 1320-1354.

- Shayo, Moses and Asaf Zussman. 2011. "Judicial Ingroup Bias in the Shadow of Terrorism." *Quarterly Journal of Economics*, 126:3, 1447-1484.
- Stalans, Loretta J. 1993. "Citizens' Crime Stereotypes, Biased Recall, and Punishment Preferences in Abstract Cases: The Educative Role of Interpersonal Sources." *Law and Human Behavior*, 17:4, 451-470.
- Starr, Sonja B. 2015. "Estimating Gender Disparities in Federal Criminal Cases." *American Law and Economics Review*, 17:1, 127-159.
- Tak, Peter J.P. 2008. *The Dutch Criminal Justice System*. Nijmegen, The Netherlands: Wolf Legal Publishers.
- Tonry, Michael. 2013. "Prosecutors and Politics in Comparative Perspective." In: Tonry, Michael (ed.), *Prosecutors and Politics. A Comparative Perspective*. Chicago, IL: University of Chicago Press, 1-33.
- Tversky, Amos and Daniel Kahneman. 1973. "Availability: A Heuristic for Judging Frequency and Probability." *Cognitive Psychology*, 5:2, 207-232.
- Van Atteveldt, Wouter, Nel Ruigrok, and Jan Kleinnijenhuis. 2005. "Associative Framing: A Unified Method for Measuring Media Frames and the Media Agenda." Unpublished manuscript.
- Van der Veer, Peter. 2006. "Pim Fortuyn, Theo van Gogh, and the Politics of Tolerance in the Netherlands." *Public Culture*, 18:1, 111-124.
- Voigt, Stefan and Alexander J. Wulf. 2019. "What Makes Prosecutors Independent? Analysing the Institutional Determinants of Prosecutorial Independence." *Journal of Institutional Economics*, 15:1, 99-120.

Table 1: Variable definitions

Variable Name	Description
<i>Outcome</i>	
Prosecute	Dummy equal to 1 if charge against suspect is prosecuted.
Drop	Dummy equal to 1 if charge against suspect is dropped.
<i>Suspect characteristics</i>	
Muslim	Dummy equal to 1 if suspect is born in country with at least 50% Muslim population.
Age [15-25]	Dummy equal to 1 if suspect's age (in years) is between 15 and 25.
Age [26-35]	Dummy equal to 1 if suspect's age (in years) is between 26 and 35.
Age [36-45]	Dummy equal to 1 if suspect's age (in years) is between 36 and 45.
Age [46-55]	Dummy equal to 1 if suspect's age (in years) is between 46 and 55.
Age \geq 56	Dummy equal to 1 if suspect's age (in years) is 56 or more.
Unemployed	Dummy equal to 1 if suspect receives unemployment benefits, welfare benefits, illness/disability benefits, other social security benefits, or has no income at all.
Retired	Dummy equal to 1 if suspect is retired.
Student	Dummy equal to 1 if suspect is a student
Employed	Dummy equal to 1 if suspect is employed or self-employed.
Criminal History	Dummy equal to 1 if suspect has been convicted for a crime in the last three years.
<i>Charge characteristics</i>	
Multiple Charges	Dummy equal to 1 if the suspect is facing multiple charges.
Abuse	Dummy equal to 1 if at least one of the charges is classified as abuse.
Threat	Dummy equal to 1 if at least one of the charges is classified as threat.
Assault	Dummy equal to 1 if at least one of the charges is classified as assault.
Crime Against Life	Dummy equal to 1 if at least one of the charges is classified as crime against life.
Weapon	Dummy equal to 1 if at least one of the charges is classified as weapon.
Other Crime	Dummy equal to 1 if at least one of the charges is classified as other crime.

Notes: The table provides variable definitions.

Table 2: Descriptive statistics

	Obs.	Mean	Std. dev.	Min.	Max.
<i>Outcome</i>					
Charged	1,193	0.79	0.41	0	1
Drop	1,193	0.04	0.18	0	1
<i>Suspect Characteristics</i>					
Muslim	1,193	0.18	0.38	0	1
Age 15-25	1,193	0.27	0.45	0	1
Age 26-35	1,193	0.30	0.46	0	1
Age 36-45	1,193	0.28	0.45	0	1
Age 46-55	1,193	0.11	0.31	0	1
Age 56+	1,193	0.04	0.20	0	1
Unemployed	1,193	0.57	0.50	0	1
Retired	1,193	0.02	0.13	0	1
Student	1,193	0.06	0.23	0	1
Employed	1,193	0.35	0.48	0	1
Criminal History	1,193	0.43	0.49	0	1
<i>Charge Characteristics</i>					
Multiple Charges	1,193	0.59	0.49	0	1
Abuse	1,193	0.47	0.50	0	1
Threat	1,193	0.22	0.41	0	1
Assault	1,193	0.03	0.16	0	1
Crime Against Life	1,193	0.13	0.34	0	1
Weapon	1,193	0.10	0.30	0	1
Other Violent Crime	1,193	0.06	0.24	0	1

Notes: The table reports the descriptive statistics for the featured variables.

Table 3: Results on tests of covariate balance

	Mean		(3) OLS Difference
	(1) Before van Gogh murder	(2) After van Gogh murder	
<i>Suspect characteristics</i>			
Muslim	0.17	0.19	0.023 (0.309) [1.000]
Age 26-35	0.30	0.29	-0.016 (0.545) [1.000]
Age 36-45	0.29	0.27	-0.019 (0.458) [1.000]
Age 46-55	0.11	0.11	0.004 (0.828) [1.000]
Age 56+	0.04	0.05	0.006 (0.615) [1.000]
Unemployed	0.58	0.56	-0.014 (0.621) [1.000]
Retired	0.01	0.02	0.010 (0.176) [1.000]
Student	0.05	0.06	0.014 (0.298) [1.000]
Criminal History	0.42	0.43	0.009 (0.764) [1.000]
<i>Charge characteristics</i>			
Multiple Charges	0.60	0.57	-0.024 (0.396) [1.000]
Abuse	0.48	0.45	-0.036 (0.213) [1.000]
Threat	0.19	0.25	0.052 (0.031) [0.496]
Assault	0.03	0.03	-0.001 (0.940) [1.000]
Crime Against Life	0.12	0.13	0.009 (0.628) [1.000]
Weapon	0.11	0.09	-0.017 (0.322) [1.000]
Other Violent Crime	0.06	0.05	-0.008 (0.575) [1.000]

Notes: The table reports the results based on tests of covariate balance. Columns (1) and (2) reports mean values for the variables respectively before and after the van Gogh murder. Column (3) reports the coefficients from an OLS regression with Post as the explanatory variable, with corresponding p -values shown in (parentheses) and p -values adjusted for multiple-hypothesis testing (Bonferroni correction) shown in [brackets].

Table 4: Main results

	Dependent variable: Prosecute						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Post Murder × Muslim	0.164 (0.086)* [0.051]*** {0.053}***	0.170 (0.086)** [0.052]*** {0.051}***	0.141 (0.086)* [0.052]** {0.045}***	0.126 (0.086) [0.062]* {0.039}***	0.134 (0.083) [0.050]** {0.051}**	0.197 (0.084)** [0.054]*** {0.054}***	0.385 (0.173)** [0.118]*** {0.125}***
Muslim	-0.130** (0.065)	-0.136** (0.065)	-0.118* (0.067)	-0.100 (0.068)	-0.098 (0.066)	-0.132** (0.063)	-0.026 (0.116)
Post Murder	-0.034 (0.032)						
Biweek FE	No	Yes	Yes	Yes	Yes	Yes	Yes
Suspect characteristics	No	No	Yes	Yes	Yes	Yes	Yes
Crime characteristics	No	No	No	Yes	Yes	Yes	Yes
Crime month-by-year FE	No	No	No	No	Yes	Yes	Yes
Interactions Post Murder and controls	No	No	No	No	No	Yes	Yes
Muslim time trend	No	No	No	No	No	No	Yes
R-squared	0.008	0.013	0.070	0.109	0.179	0.217	0.219
Observations	1,193	1,193	1,193	1,193	1,193	1,193	1,193

Notes: The table reports OLS results. The dependent variable is Prosecute. The unit of observation is a charge. The sample includes all charges processed within 6 weeks before and 6 weeks after the van Gogh murder (12-week bandwidth). Respective standard errors reported in (parentheses), [square brackets], and {curly brackets} are clustered at the level of the case, city of crime, and accused's country of birth. ***, ** and *, which are for the Post Murder×Muslim coefficient estimate only shown next to the applicable standard errors (as opposed to next to the coefficient estimate), denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 5: Results based on a dynamic specification

	Dependent variable: Prosecute		
	(1)	(2)	(3)
3rd Biweek Before Murder × Muslim	-0.031 (0.154)	0.006 (0.150)	-0.008 (0.144)
2nd Biweek Before Murder × Muslim	0.048 (0.173)	0.057 (0.168)	0.069 (0.155)
1st Biweek After Murder × Muslim	0.267* (0.142)	0.254* (0.137)	0.323* (0.142)
2nd Biweek After Murder × Muslim	0.198 (0.133)	0.223* (0.132)	0.268** (0.132)
3rd Biweek After Murder × Muslim	-0.036 (0.157)	0.013 (0.159)	0.079 (0.156)
Muslim	-0.101 (0.116)	-0.117 (0.114)	-0.149 (0.112)
Biweek FE	Yes	Yes	Yes
Suspect characteristics	Yes	Yes	Yes
Crime characteristics	Yes	Yes	Yes
Crime month-by-year FE	No	Yes	Yes
Interactions Post Murder and controls	No	No	Yes
R-squared	0.116	0.184	0.222
Observations	1,193	1,193	1,193

Notes: The table reports OLS results. The dependent variable is Prosecute. The unit of observation is a charge. The sample includes all charges processed within 6 weeks before and 6 weeks after the van Gogh murder (12-week bandwidth). Standard errors reported in parentheses are clustered at the case level. ***, ** and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 6: Results when varying the bandwidth, static model

	Dependent variable: Prosecute				
	Bandwidth				
	(1) 8 weeks	(2) 10 weeks	(3) 14 weeks	(4) 16 weeks	(5) 20 weeks
Post Murder × Muslim	0.198* (0.101)	0.278*** (0.088)	0.142* (0.076)	0.115 (0.070)	0.056 (0.064)
Muslim	Yes	Yes	Yes	Yes	Yes
Biweek FE	Yes	Yes	Yes	Yes	Yes
Suspect characteristics	Yes	Yes	Yes	Yes	Yes
Crime characteristics	Yes	Yes	Yes	Yes	Yes
Crime month-by-year FE	Yes	Yes	Yes	Yes	Yes
Interactions Post Murder and controls	Yes	Yes	Yes	Yes	Yes
R-squared	0.266	0.276	0.205	0.196	0.167
Observations	747	933	1,401	1,575	1,958

Notes: The table reports OLS results. The dependent variable is Prosecute. The unit of observation is a charge. Each column shows results for a different bandwidth. Bandwidth of length w weeks implies that the applicable sample includes all charges processed within $w/2$ weeks before and $w/2$ weeks after the van Gogh murder. Standard errors reported in parentheses are clustered at the case level. ***, ** and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 7: Results when varying the bandwidth, dynamic model

	Dependent variable: Prosecute		
	Bandwidth		
	(1) 8 weeks	(2) 16 weeks	(3) 20 weeks
5th Biweek Before Murder × Muslim			0.201 (0.148)
4th Biweek Before Murder × Muslim		0.182 (0.134)	0.166 (0.139)
3rd Biweek Before Murder × Muslim		0.019 (0.143)	0.040 (0.149)
2nd Biweek Before Murder × Muslim	0.159 (0.152)	0.056 (0.159)	0.043 (0.165)
1st Biweek After Murder × Muslim	0.263* (0.148)	0.316** (0.139)	0.312** (0.141)
2nd Biweek After Murder × Muslim	0.281** (0.141)	0.241* (0.128)	0.247* (0.130)
3rd Biweek After Murder × Muslim		0.058 (0.156)	0.032 (0.157)
4th Biweek After Murder × Muslim		0.118 (0.131)	0.109 (0.132)
5th Biweek After Murder × Muslim			0.070 (0.146)
Muslim	Yes	Yes	Yes
Biweek FE	Yes	Yes	Yes
Suspect characteristics	Yes	Yes	Yes
Crime characteristics	Yes	Yes	Yes
Crime month-by-year FE	Yes	Yes	Yes
Interactions Post Murder and controls	Yes	Yes	Yes
R-squared	0.268	0.202	0.174
Observations	747	1,575	1,958

Notes: The table reports OLS results. The dependent variable is Prosecute. The unit of observation is a charge. Each column shows results for a different bandwidth. Bandwidth of length w weeks implies that the applicable sample includes all charges processed within $w/2$ weeks before and $w/2$ weeks after the van Gogh murder. Standard errors reported in parentheses are clustered at the case level. ***, ** and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 8: Results for placebo treatment dates

	Dependent variable: Prosecute		
	(1)	(2)	(3)
Post Murder × Muslim	-0.072 (0.055)	0.003 (0.057)	-0.067 (0.064)
Muslim	Yes	Yes	Yes
Biweek FE	Yes	Yes	Yes
Suspect characteristics	Yes	Yes	Yes
Crime characteristics	Yes	Yes	Yes
Crime month-by-year FE	Yes	Yes	Yes
Interactions Post Murder and controls	Yes	Yes	Yes
R-squared	0.344	0.393	0.287
Observations	913	994	1,052

Notes: The table reports OLS results. The dependent variable is Prosecute. The unit of observation is a charge. Column (1) shows results for cases processed within six weeks before or after November 2, 2001. Column (2) shows results for cases processed within six weeks before or after November 2, 2002. Column (3) shows results for cases processed within six weeks before or after November 2, 2003. Standard errors reported in parentheses are clustered at the case level.

Table 9: Results when varying definition of treatment and control group

	Dependent variable: Prosecute						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Post Murder × Muslim	0.197** (0.084)	0.197** (0.084)	0.191** (0.084)	0.177** (0.085)	0.136** (0.067)	0.110 (0.116)	0.208** (0.088)
Muslim	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Biweek FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Suspect characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Crime characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Crime month-by-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Interactions Post Murder and controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.219	0.219	0.219	0.217	0.216	0.216	0.233
Observations	1,193	1,193	1,193	1,193	1,193	1,193	910

Notes: The table reports OLS results. The dependent variable is Prosecute. The unit of observation is a charge. The sample includes all charges processed within 6 weeks before and 6 weeks after the van Gogh murder (12-week bandwidth). Columns (1) to (4) show results based on varying the definition of a Muslim-majority country based on which we define the treated group. We increase the threshold value of the proportion of population that self-identifies as Muslim in the suspect's country of origin from the baseline 50 percent to 60 percent (column (1)), 70 percent (column (2)), 80 percent (column (3)) and 90 percent (column (4)). Column (5) shows results when the treated group includes not only suspects who were themselves born in a Muslim-majority country but also suspects whose parents had been born in a Muslim-majority country. Column (6) shows results when the treated group includes suspects whose country of birth is Morocco. Column (7) shows results when we re-define our control group to consist only of suspects who were not born abroad and whose parents were also not born abroad. Standard errors reported in parentheses are clustered at the case level. ***, ** and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 10: Results based on the sample of crimes committed before van Gogh murder

	Dependent variable: Prosecute					
	(1)	(2)	(3)	(4)	(5)	(6)
Post Murder × Muslim	0.200*	0.175*	0.157	0.174*	0.212**	
	(0.105)	(0.103)	(0.102)	(0.103)	(0.105)	
3rd Biweek Before Murder × Muslim						-0.005
						(0.145)
2nd Biweek Before Murder × Muslim						0.073
						(0.156)
1st Biweek After Murder × Muslim						0.326**
						(0.165)
2nd Biweek After Murder × Muslim						0.303*
						(0.175)
3rd Biweek After Murder × Muslim						0.084
						(0.194)
Biweek FE	Yes	Yes	Yes	Yes	Yes	Yes
Suspect characteristics	No	Yes	Yes	Yes	Yes	Yes
Crime characteristics	No	No	Yes	Yes	Yes	Yes
Crime month-by-year FE	No	No	No	Yes	Yes	Yes
Interactions Post Murder and controls	No	No	No	No	Yes	Yes
R-squared	0.022	0.083	0.134	0.202	0.250	0.254
Observations	947	947	947	947	947	947

Notes: The table reports OLS results. The dependent variable is Prosecute. The unit of observation is a charge. Columns (1) to (6) show results based on specifications as featured in columns (2) to (6) of Table 4 and column (3) of Table 5, respectively, but when restricting the estimating sample to crimes committed prior to van Gogh murder. Standard errors reported in parentheses are clustered at the case level. ***, ** and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 11: Quantifying sensitivity to omitted variable bias

Part A: Abs. value of δ for $\beta=0$		
(1) $R_{max}=\min\{1.3R^2,1\}$	(2) $R_{max}=\min\{2R^2,1\}$	(3) $R_{max}=\min\{4R^2,1\}$
4.33	4.21	0.49
Part B: Bounds on estimate of β for $\delta=1$		
(4) $R_{max}=\min\{1.3R^2,1\}$	(5) $R_{max}=\min\{2R^2,1\}$	(6) $R_{max}=\min\{4R^2,1\}$
[0.198, 0.414]	[0.198, 2.779]	[0.198, 13.466]

Notes: The table summarizes the results of Oster (2019) tests of the sensitivity to omitted variable bias of the estimates of the focal difference-in-difference coefficient on Post Murder \times Muslim based on the specification reported in column (6) of Table 4. Part A of the table shows the value of δ , the ratio of the importance of selection on unobservables versus observables for which the effect of interest would be equal to zero. Part B shows the bounds for the estimated effects when $\delta=1$. Columns (1) and (4) show results based on Oster-recommended value of R_{max} , the value of R^2 from a hypothetical regression containing all relevant (observed and unobserved) controls. Columns (2) and (5), as well as (3) and (6), show results based on even more conservative scenarios.

Table 12: Results for non-violent crimes

	Dependent variable: Prosecute				
	Non-violent crime type:				
	(1) Property	(2) Drugs	(3) Destruction	(4) Traffic	(5) Other
Post Murder × Muslim	-0.005 (0.042)	0.171 (0.118)	-0.118 (0.100)	-0.255*** (0.083)	0.063 (0.075)
Muslim	Yes	Yes	Yes	Yes	Yes
Biweek FE	Yes	Yes	Yes	Yes	Yes
Suspect characteristics	Yes	Yes	Yes	Yes	Yes
Crime characteristics	Yes	Yes	Yes	Yes	Yes
Crime month-by-year FE	Yes	Yes	Yes	Yes	Yes
Interactions Post Murder and controls	Yes	Yes	Yes	Yes	Yes
R-squared	0.213	0.242	0.294	0.243	0.402
Observations	1,826	484	552	1,069	597

Notes: The table reports OLS results. The dependent variable is Prosecute. The unit of observation is a charge. The sample includes all charges processed within 6 weeks before and 6 weeks after the van Gogh murder (12-week bandwidth). The table reports the results for different types of non-violent crimes: property crimes (column (1)), drug crimes ((column (2)), destruction crimes (column (3)), traffic crimes (column (4)) and other crimes (column (5)). Standard errors reported in parentheses are clustered at the case level. ***, ** and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 13: Results for crimes processed by non-Amsterdam prosecution offices

	Dependent variable: Prosecute								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post Murder × Muslim	0.027 (0.049)	-0.072 (0.061)	0.021 (0.101)	-0.043 (0.101)	-0.064 (0.095)	-0.018 (0.059)	-0.124 (0.100)	0.081 (0.115)	0.029 (0.102)
Muslim	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Biweek FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Suspect characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Crime characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Crime month-by-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Interactions Post Murder and controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.245	0.251	0.292	0.180	0.322	0.206	0.225	0.243	0.241
Observations	1,561	1,493	805	1,295	1,016	2,461	848	751	992

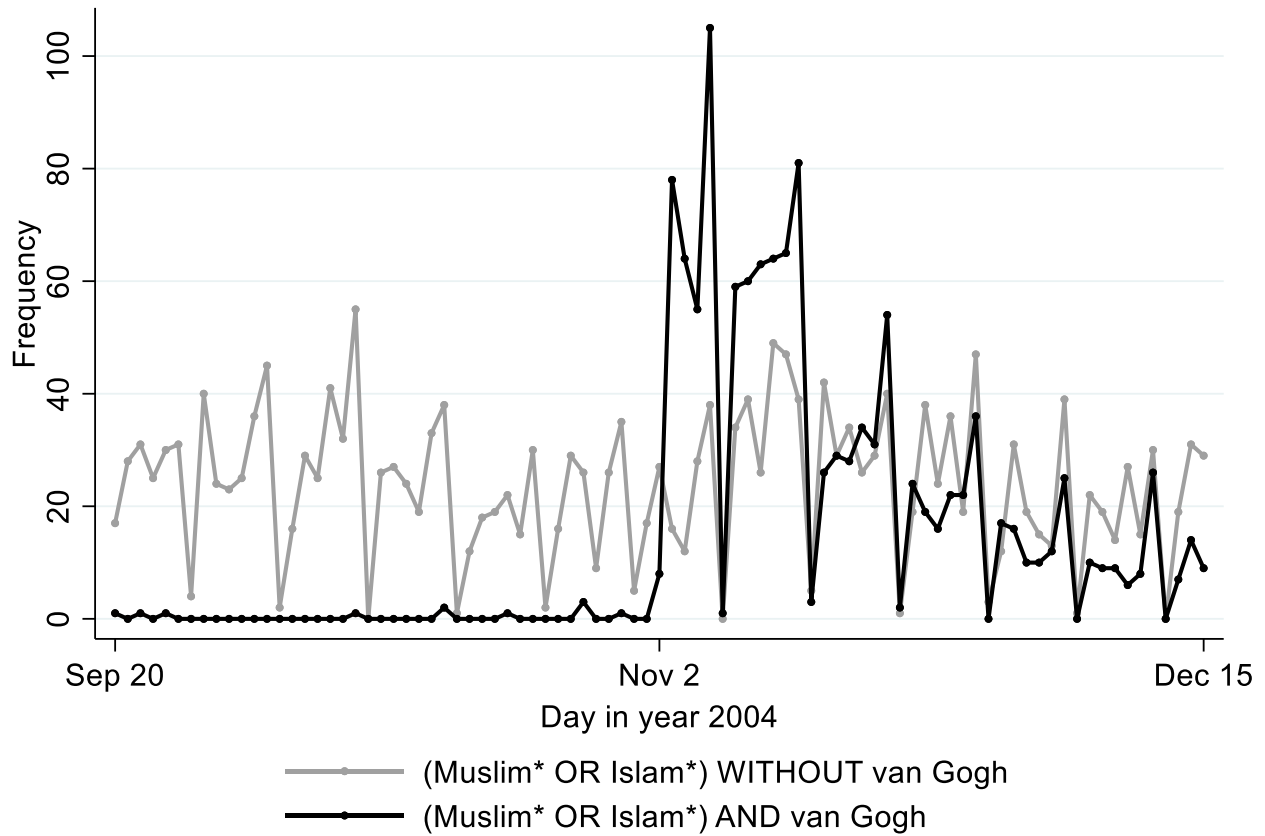
Notes: The table reports OLS results. The dependent variable is Prosecute. The unit of observation is a charge. The samples include all charge processed within 6 weeks before and 6 weeks after the van Gogh murder (12-week bandwidth). Each column shows results for charges processed by a different non-Amsterdam prosecution office in the Netherlands. Standard errors reported in parentheses are clustered at the case level. ***, ** and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 14: Results using an alternative outcome variable

	Dependent variable: Drop				
	(1)	(2)	(3)	(4)	(5)
Post Murder × Muslim	-0.056 (0.040)	-0.052 (0.038)	-0.059* (0.035)	-0.066** (0.033)	-0.069** (0.031)
Muslim	Yes	Yes	Yes	Yes	Yes
Biweek FE	Yes	Yes	Yes	Yes	Yes
Suspect characteristics	No	Yes	Yes	Yes	Yes
Crime characteristics	No	No	Yes	Yes	Yes
Crime month-by-year FE	No	No	No	Yes	Yes
Interactions Post Murder and controls	No	No	No	No	Yes
R-squared	0.009	0.013	0.079	0.128	0.147
Observations	1,193	1,193	1,193	1,193	1,193

Notes: The table reports OLS results. The dependent variable is Drop. The unit of observation is a charge. The sample includes all charges processed within 6 weeks before and 6 weeks after the van Gogh murder (12-week bandwidth). Columns (1) to (5) show results based on specification analogous to those featured in columns (2) to (6) of Table 4, respectively, upon replacing the binary indicator Prosecute with a binary indicator Drop equal to one if the charge in the given criminal case against a particular suspect processed during the given time period was altogether dropped, and zero otherwise. Standard errors reported in parentheses are clustered at the case level. ***, ** and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Figure 1: Newspaper coverage of events



Notes: The figure plots the time series of the daily incidence of different accounts (articles, editorials, and other features) in eight major Dutch newspapers shortly before and after the Van Gogh murder. The series in black shows the temporal evolution of the daily frequency of accounts that contain words 'Muslim' or 'Islam' (and their derivative grammatical forms in Dutch) as well as 'van Gogh'. The series in gray shows the temporal evolution of the daily frequency of accounts containing words 'Muslim' or 'Islam' that do not include 'van Gogh'.

Figure 2: Parallel time trends

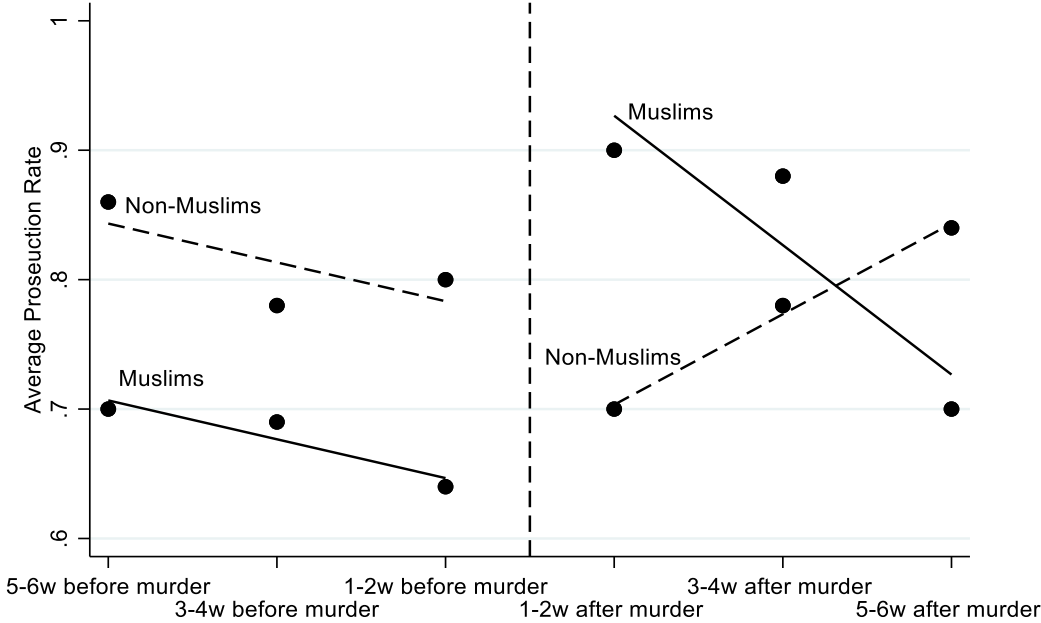



Figure 3: Summary of dynamic effects



Notes: The figure summarize the effects based on the estimates in Table 5.

Figure 4: Summary of results from randomly re-assigning treated group status

