ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

DOTTORATO DI RICERCA IN

ECONOMICS

Ciclo 36

Settore Concorsuale: 13/A1 - ECONOMIA POLITICA

Settore Scientifico Disciplinare: SECS-P/01 - ECONOMIA POLITICA

ESSAYS IN POLITICAL ECONOMY AND CRIME ECONOMICS

Presentata da: Marco Rosso

Coordinatore Dottorato

Matteo Barigozzi

Supervisore

Paolo Vanin

Esame finale anno 2026

To Paolo, for his constant guidance and care,
both throughout my doctoral journey and beyond.

To Chiara (and Caterina),
for showing me the place where I truly belong.

To Clara, Walter, and Stefano,
for 34 years of support and patience.

And to all those who have walked this path with me,
whether briefly or at length.

Abstract

This dissertation studies how non-monetary frictions—such as temptation, information salience, and time constraints—shape individual behavior in contexts where formal institutions are otherwise well-functioning. Combining theoretical analysis with micro-level empirical evidence, the three chapters examine how these frictions affect selection into public employment and participation in democratic processes.

The *first chapter* develops a behavioral theory of occupational selection in the public sector. Introducing self-control costs into a standard model of career choice, the analysis shows that corruption generates a dual selection effect: while illicit rents attract low-motivation individuals, the psychological costs of resisting temptation deter highly motivated agents. The model identifies an institutional tipping point at which corruption switches from expanding public employment while degrading workforce quality to acting as a severe screening device that improves average quality but reduces participation. These results clarify why similar anti-corruption environments can produce sharply different selection outcomes across countries.

The *second chapter* examines how crime-related information affects individual voting behavior. Using geolocated data that link retrospective voting choices to local crime news coverage in Bologna across multiple national and municipal elections, the analysis exploits within-individual variation in exposure to nearby crime reports. The results show that aggregate crime salience has weak and unstable effects, while crime attributed to immigrants generates systematic electoral responses: voters shift away from parties with ambiguous positions on immigration toward parties emphasizing law and order. In local elections, immigrant-related crime increases abstention, whereas crimes committed by natives lead to punishment of incumbents. These findings highlight the role of identity-based framing in mediating the political impact of crime.

The *third chapter* studies how parenthood and parental age shape electoral participation. Using administrative data covering the universe of registered voters in Bologna, the analysis follows individuals over time as their family circumstances evolve. Once permanent individual heterogeneity is accounted for, parenthood is not associated with lower turnout on average. However, substantial life-cycle heterogeneity emerges: parents of infants and preschoolers—especially mothers—exhibit sizable turnout penalties at younger ages, which decline steadily with parental age and disappear by around age forty. These results indicate that periods of intensive childcare impose temporary opportunity costs on political participation.

Taken together, the three chapters show how behavioral frictions and life-cycle constraints can generate distortions in selection and participation even in settings with low formal barriers to entry and participation. By emphasizing micro-level mechanisms rather than institutional failures alone, the dissertation contributes to the understanding of public-sector composition, electoral behavior, and democratic representation.

Contents

The Tipping Point of Temptation: Occupational Selection and Integrity in the Public Sector*

Marco Rosso†

Abstract

This paper develops a behavioral theory of occupational selection based on endogenous self-control costs, addressing the long-standing empirical ambiguity regarding the quality of the public sector workforce in environments characterized by corruption and moral frictions. The framework integrates self-control costs and temptation into a standard model of occupational choice, drawing on the utility framework of Gul and Pesendorfer (2001). We show that intrinsically motivated (honest) agents face disproportionately higher psychological costs when resisting temptation, generating a dual selection effect: low-motivation types are attracted to public employment, while highly motivated types are increasingly deterred. To discipline these opposing forces, the analysis establishes three general principles governing institutional selection, supported by analytical derivations and general selection arguments under weak regularity conditions. These principles are shown to extend beyond the benchmark environment and to hold under weak regularity conditions on the joint distribution of ability and honesty. First, we identify a critical institutional tipping point, λ^*, that determines the selection regime: below it, corruption deteriorates workforce quality (*"more but worse"*); above it, corruption acts as a severe screening device, improving average quality (*"less but better"*). Second, we show that selection outcomes are fundamentally conditional on the societal correlation between ability and honesty. Third, the model provides a novel rationale for high public-sector wages, demonstrating that sufficiently large salaries attenuate the selective power of corruption by shielding high-motivation agents from self-control costs. Overall, the paper clarifies the mechanisms shaping workforce composition in morally frictional environments and contributes to the literature on occupational selection, public service motivation, and institutional design.

Keywords: occupational selection; self-control and temptation; corruption; public sector labor markets; institutional design.
JEL: D73; J45; D90; H83.

*I thank participants in seminars and conferences in Bologna for helpful comments. I am particularly grateful to Tommy E. Murphy, Juan F. Vargas, and Paolo Vanin for their valuable feedback. I gratefully acknowledge financial support from Fondazione Carisbo. All errors are my own. The pronoun "we" is used throughout the paper for convenience.

†Department of Economics, University of Bologna (email: marco.rosso4@unibo.it).

1. Introduction

A central challenge in public economics is understanding how institutional environments shape the selection and retention of high-quality workers in the public sector. Standard models of occupational choice emphasize monetary incentiveswages, monitoring, and sanctionsand predict that improving these incentives should uniformly raise workforce quality (Becker, 1968; Ehrlich, 1973). Yet empirical evidence on this prediction is mixed and highly context-dependent.

This ambiguity is particularly evident in environments characterized by corruption. While corruption is widely recognized as welfare-reducing, its effect on the composition of the public-sector workforce remains theoretically unsettled and empirically debated. In some contexts, corruption appears to attract individuals with low intrinsic motivation or weaker ethical standards, as documented for India by Hanna and Wang (2017). In others, such as Denmark, public-sector employment is strongly associated with pro-social motivation and integrity (Barfort et al., 2019). Related evidence suggests that corruption opportunities may simultaneously attract dishonest individuals and deter honest ones, potentially due to the psychological or moral costs of operating in corrupt environments (Brassiolo et al., 2021; Konrad et al., 2021). Taken together, these findings highlight the absence of a unified theoretical framework capable of reconciling the heterogeneous effects of corruption on public-sector selection.

This paper addresses this gap by introducing a behavioral frictionthe cost of self-controlinto an otherwise standard model of occupational choice. Drawing on the literature on temptation and self-control (Gul and Pesendorfer, 2001; Bénabou and Tirole, 2011), we model corruption as a short-run temptation that generates a psychological cost for agents who resist it. Crucially, this cost is increasing in intrinsic motivation (honesty), implying that high-integrity agents face a disproportionately higher burden when operating in corrupt institutions. As a result, corruption does not merely distort incentives through rents and sanctions, but also imposes a non-pecuniary "moral tax" on public service.

The analysis yields three central insights. First, in a tractable benchmark environment, corruption can affect public-sector participation in a non-monotonic way. When institutional temptation is limited, corruption expands public employment while deteriorating average workforce quality; when temptation is sufficiently severe, corruption instead acts as a selective deterrent, reducing participation but improving composition. Second, selection outcomes depend critically on the societal correlation between ability and honesty. When these traits are negatively correlated, policymakers face an unavoidable trade-off: policies that attract competence tend to deter integrity, and vice versa. Third, the analysis provides a novel interpretation of high public-sector wages. Beyond attracting skills, higher wages attenuate the self-control costs associated with corruption, thereby shielding intrinsically motivated agents from adverse selection.

These mechanisms help organize several stylized facts in the empirical literature. In high-corruption environments such as India, pro-social preferences are negatively associated with public-sector career choices (Hanna and Wang, 2017), consistent with a setting in which ability and honesty are misaligned. In contrast, low-corruption environments such as Denmark display a positive association between pro-social motivation and public employment (Barfort et al., 2019). Experimental evidence further shows that rent opportunities reduce honest take-up and shift the composition of entrants toward dishonest types (Brassiolo et al., 2021), in line with the selective forces highlighted by the model.

The paper relates to several strands of the literature. First, it contributes to work on occupational sorting and public-sector labor markets (Roy, 1951; Lazear, 1979; Delfgaauw and Dur, 2010; Barigozzi et al., 2018; Besley, 2004; Banuri and Keefer, 2016) by introducing an institutionally induced, non-monetary cost that shapes selection. Second, it complements studies on corruption and public-sector performance (Caselli and Morelli, 2004; Acemoglu and Verdier, 2000; Dal Bó and Di Tella, 2003; Dal Bó et al., 2013; Acemoglu et al., 2013) by providing a behavioral microfoundation for the self-exclusion of high-integrity agents from corrupt institutions. Finally, it connects insights from the literature on values, norms, and political selection (Corneo and Jeanne, 2009, 2010; Bernheim and Kartik, 2014) to models of institutional design, highlighting the role of moral frictions in workforce composition.

The remainder of the paper proceeds as follows. Section 2 introduces the model; Section 3 analyzes occupational selection; Section 4 presents benchmark results under uniform distributions; Section 6 studies institutional design and the endogenization of temptation. Section 5 develops a general selection result under arbitrary distributions, establishing the existence and robustness of the institutional tipping point; Section 7 concludes. Formal proofs are provided in Appendix A, while extensions and robustness results based on monotone-set inclusion and moving selection frontiers are collected in Appendix B.

2. Model

2.1. Environment and Preferences

We consider an economy with two sectors, a public sector (P) and a private sector (p). Workers are heterogeneous along two dimensions: ability θ, which determines productivity in the private sector, and intrinsic motivation γ, which captures the non-pecuniary value of public-sector employment. Agents freely choose their sector of employment.

In the private sector, workers receive a payoff increasing in ability:

$$u_p = \theta. \tag{1}$$

In the public sector, workers can behave either honestly (h) or corruptly (c). We adopt a linear and additively separable utility function.[1] Expected utility is given by

$$u_{\mathrm{P}}(h) = w + \gamma \tag{2}$$

under honest behavior, and by

$$u_{\mathrm{P}}(c) = w + b \tag{3}$$

under corrupt behavior. Public employment provides a fixed wage w and a non-monetary premium from honesty, while corruption yields an illicit gain b. As in Dal Bó et al. (2006), corruption opportunities are confined to the public sector.

2.2. Temptation and Self-Control

Public employment exposes workers to corruption opportunities that generate short-run temptation. We model this tension using the dual-self framework of Gul and Pesendorfer (2001) and Cervellati and Vanin (2013).

Temptation arises from an overestimation of the immediate gains from corruption. This is captured by two rankings: a commitment utility u_{P}, defined above, and a temptation utility v_{P}. The two rankings differ only in the valuation of corrupt behavior. Specifically, temptation inflates the perceived return from corruption by a factor $\lambda > 1$, so that

$$v_{\mathrm{P}}(h) = w + \gamma, \tag{4}$$

and

$$v_{\mathrm{P}}(c) = w + \lambda b. \tag{5}$$

Under temptation, corruption maximizes the temptation ranking if and only if $\lambda b > \gamma$. Self-control problems arise whenever the temptation ranking and the commitment ranking disagree. Following Gul and Pesendorfer (2001) and Cervellati and Vanin (2013), a choice is *tempting* if it maximizes v_{P} but not u_{P}. Resisting temptation requires exerting self-control, the cost of which equals the forgone temptation utility. Formally, choosing an action $x \in \{c, h\}$ that does not maximize v_{P} entails a

[1] The main results do not depend on linearity; this formulation is chosen for expositional simplicity.

self-control cost

$$\max_{y\in\{c,h\}} \Big[v_{\mathrm{P}}(y) - v_{\mathrm{P}}(x)\Big].$$

Individuals choose the action that maximizes commitment utility net of self-control costs. Overall preferences are therefore represented by

$$U(x) = \max_{x\in\{c,h\}} \left\{ u_{\mathrm{P}}(x) - \left[\max_{y\in\{c,h\}} v_{\mathrm{P}}(y) - v_{\mathrm{P}}(x) \right] \right\}. \tag{6}$$

We first characterize optimal behavior within the public sector and then study self-selection across sectors.

2.3. Optimal Behavior in the Public Sector

In the absence of self-control problems, corruption is chosen whenever $\gamma < b$, while honesty is chosen whenever $\gamma > \lambda b$. For intermediate levels of intrinsic motivation, $\gamma \in [b, \lambda b]$, honest behavior maximizes commitment utility but corruption is tempting. We refer to this interval as the *temptation range* (Figure 1).

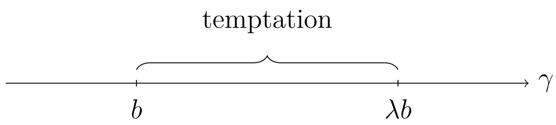

Fig. 1: The temptation range in the public sector by intrinsic motivation.

Within the temptation range, the optimal choice depends on the comparison between the utility loss from yielding to temptation, $\gamma - b$, and the cost of exerting self-control, $\lambda b - \gamma$. The two costs are equal at $\gamma = \frac{1+\lambda}{2}b$. Agents with lower motivation give in to temptation, while sufficiently motivated agents resist and remain honest (Figure 2).

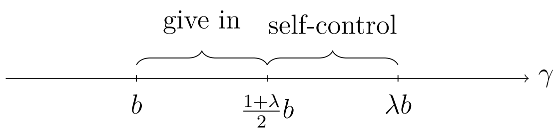

Fig. 2: Give-in and self-control regions within the temptation range.

The optimal behavior in the public sector is summarized in the following lemma.

Lemma 1. Optimal choices in the public sector: *For any* $\{\gamma, b, \lambda\}$ *with* $b > 0$ *and* $\lambda > 1$*:*

i) If $\gamma < b$*, corruption is optimal under both commitment and temptation.*

ii) If $\gamma \in [b, \lambda b]$*, honesty maximizes commitment utility but corruption is tempting:*

a) *if $\gamma < \frac{1+\lambda}{2}b$, agents give in to temptation and choose corruption;*

b) *if $\gamma > \frac{1+\lambda}{2}b$, agents exert self-control and remain honest.*

iii) *If $\gamma > \lambda b$, honesty is optimal under both rankings.*

Intuition. Temptation affects only agents with intermediate intrinsic motivation. Low-γ agents find corruption optimal even without self-control problems, while high-γ agents are never tempted. Self-control costs therefore distort behavior precisely among intrinsically motivated workers.

Lemma 1 implies a partition of motivation into a *corruption range*, $\gamma \in [0, \frac{1+\lambda}{2}b)$, and an *honesty range*, $\gamma \in (\frac{1+\lambda}{2}b, \infty)$ (Figure 3).

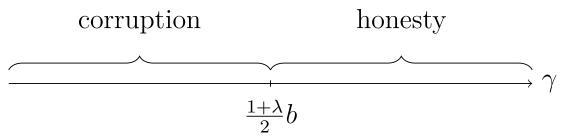

Fig. 3: Corruption and honesty regions by intrinsic motivation.

Expected utility in the public sector is therefore given by

$$U(\gamma) = \begin{cases} w + b, & \gamma < \frac{1+\lambda}{2}b, \\ w + 2\gamma - \lambda b, & \gamma \in \left(\frac{1+\lambda}{2}b, \lambda b\right], \\ w + \gamma, & \gamma > \lambda b. \end{cases}$$

Lemma 2. Utility in the public sector: *For any $\{\gamma, b, \lambda\}$ with $b > 0$ and $\lambda > 1$:*

i) *if $\gamma < \frac{1+\lambda}{2}b$, then $U(c) = w + b$;*

ii) *if $\gamma \in (\frac{1+\lambda}{2}b, \lambda b]$, then $U(h) = w + 2\gamma - \lambda b$;*

iii) *if $\gamma > \lambda b$, then $U(h) = w + \gamma$.*

Self-control versus reputation. A natural point of comparison is with models in which integrity is sustained by reputational or signaling concerns (Bénabou and Tirole, 2006; Bénabou and Tirole, 2011). In such frameworks, dishonest behavior is deterred by the anticipated reaction of others and therefore hinges on observability, beliefs, and repeated interaction. The mechanism studied here is conceptually distinct. Self-control costs operate through an *internal conflict* and arise even when actions are unobserved and perfectly anticipated. As a result, reputation primarily disciplines behavior *conditional on entry*, whereas self-control shapes *occupational selection itself* by altering the relative attractiveness of public employment ex ante.

The two mechanisms are therefore complementary rather than competing. In environments where reputational incentives are weak or absent—for instance, when monitoring is imperfect, careers are short, or corruption is pervasive—self-control costs remain operative and can generate systematic self-selection of high-integrity agents out of the public sector. This distinction clarifies why corruption may deter honest workers even when reputational sanctions are limited, and why policies targeting monitoring and transparency need not neutralize the selection effects emphasized in this paper.

Empirical interpretation. To guide empirical interpretation, model primitives map naturally to observable proxies: intrinsic honesty (γ) to pro-social behavior (e.g. dictator-game giving), ability (θ) to exam scores or cognitive tests, corruption rents (b) to bribe exposure or illicit gain opportunities, and temptation intensity (λ) to institutional features such as monitoring intensity or payoff immediacy. These mappings are used throughout the paper to interpret the theoretical results and to connect them to the existing empirical literature.

3. Self-selection

This section characterizes how corruption opportunities and temptation shape (i) the size of the public sector and (ii) the composition of its workforce. We proceed in steps, starting from two benchmarks and then returning to the general model.

3.1. Two benchmarks

We consider two benchmark environments.

i) **No corruption opportunities:** $b = 0$. Public employment entails no illicit gain, hence no conflict between honesty and corruption.

ii) **Corruption without temptation:** $b > 0$ and $\lambda = 1$. Corruption opportunities exist, but the temptation and commitment rankings coincide, so agents face no self-control problem.

3.2. Selection without corruption ($b = 0$)

If $b = 0$, public-sector utility is $u_P(h) = w + \gamma$, while private-sector utility is $u_p = \theta$. An agent selects into the public sector if and only if

$$w + \gamma \geq \theta. \tag{7}$$

Figure 4 depicts the participation region. The frontier $\theta = w + \gamma$ separates agents who enter the public sector from those who prefer the private sector.

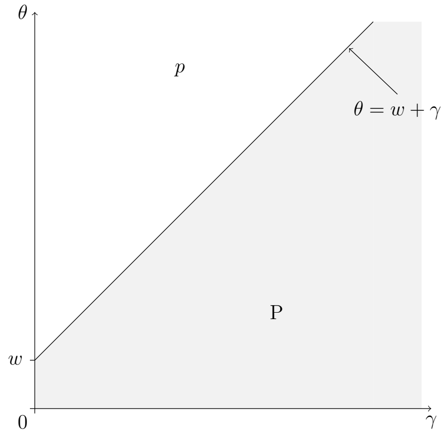

Fig. 4: Selection into the public sector in the absence of corruption ($b = 0$). *Note:* The gray region collects types (θ, γ) such that $w + \gamma \geq \theta$.

Proposition 1. *In the absence of corruption opportunities ($b = 0$), an increase in the public wage w:*

i) increases the share of workers selecting into the public sector;

ii) has an ambiguous effect on average ability and intrinsic motivation among public-sector entrants.

The first statement follows directly from the outward shift of the frontier $\theta = w + \gamma$. The second statement is inherently compositional: the effect on average ability and motivation depends on the joint distribution of (θ, γ) (and in particular on their dependence structure), hence it can differ across settings.

3.3. Corruption without temptation ($b > 0$, $\lambda = 1$)

Let $b > 0$ and $\lambda = 1$. In this case the temptation and commitment rankings coincide and behavior within the public sector is standard: agents choose honesty if and only if $\gamma \geq b$.

Selection now depends on the relevant public-sector payoff. Agents with $\gamma < b$ anticipate behaving corruptly and thus compare $u_P(c) = w + b$ to $u_p = \theta$. Agents with $\gamma \geq b$ anticipate behaving honestly and compare $u_P(h) = w + \gamma$ to $u_p = \theta$. Figure 5 summarizes the participation region.

Relative to the no-corruption benchmark, introducing $b > 0$ has two effects. First, it changes behavior for low-motivation public-sector workers (those with $\gamma < b$) who, once in the public sector, choose corruption. Second, it enlarges the participation set by attracting additional low-γ agents from the private sector (the patterned region in Figure 5).

Proposition 2. *In the absence of temptation ($\lambda = 1$), an increase in corruption opportunities b:*

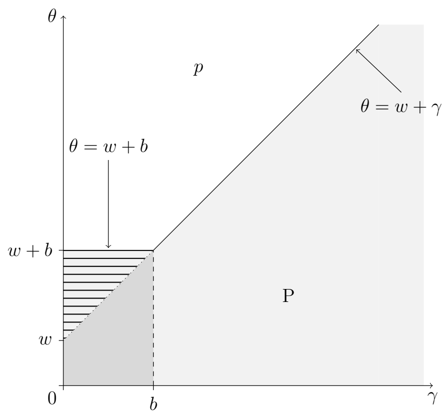

Fig. 5: Selection under corruption without temptation ($b > 0$, $\lambda = 1$). *Note:* For $\gamma < b$, public employment yields $w + b$ and the cutoff is $\theta = w + b$; for $\gamma \geq b$, public employment yields $w + \gamma$ and the cutoff is $\theta = w + \gamma$.

i) increases the share of workers selecting into the public sector;

ii) increases the share of corrupt workers in the public sector;

iii) has an ambiguous effect on average ability and intrinsic motivation among public-sector entrants.

3.4. Corruption with temptation ($b > 0$, $\lambda > 1$)

We now return to the general model. By Lemma 2, the relevant public-sector payoff depends on γ:

$$U_P(\gamma) = \begin{cases} w + b, & \gamma < \frac{1+\lambda}{2}b, \\ w + 2\gamma - \lambda b, & \gamma \in \left(\frac{1+\lambda}{2}b, \lambda b\right], \\ w + \gamma, & \gamma > \lambda b. \end{cases}$$

Selection compares $U_P(\gamma)$ to $u_p = \theta$. This yields three frontiers: (i) for $\gamma < \frac{1+\lambda}{2}b$, the cutoff is $\theta = w + b$; (ii) for $\gamma \in (\frac{1+\lambda}{2}b, \lambda b]$, the cutoff is $\theta = w + 2\gamma - \lambda b$; (iii) for $\gamma > \lambda b$, the cutoff is $\theta = w + \gamma$. Figure 6 depicts the implied participation set.

Temptation matters only in the intermediate region $\gamma \in [b, \lambda b]$, where agents incur either a commitment loss from giving in or a self-control cost from resisting. This generates three forces relative to the $\lambda = 1$ case: (i) some low-γ public employees switch behavior and become corrupt (giving in region); (ii) some intermediate-γ public employees remain honest but at lower utility (self-control region); (iii) most importantly, the reduction in $U_P(\gamma)$ for intermediate γ shrinks the set of types who select into the public sector.

To illustrate the participation effect, Figure 7 highlights the region that exits the public sector once corruption becomes tempting.

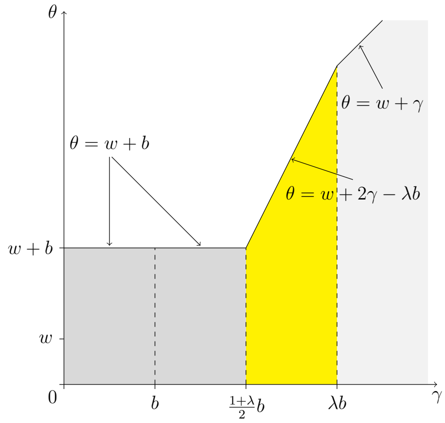

Fig. 6: Selection under tempting corruption ($b > 0$, $\lambda > 1$). *Note:* The relevant participation cutoff is $\theta = w + b$ for $\gamma < \frac{1+\lambda}{2}b$, $\theta = w + 2\gamma - \lambda b$ for $\gamma \in (\frac{1+\lambda}{2}b, \lambda b]$, and $\theta = w + \gamma$ for $\gamma > \lambda b$.

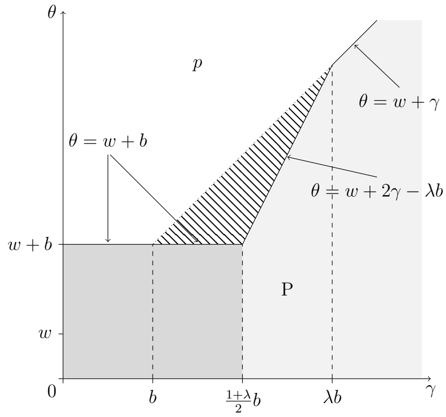

Fig. 7: The participation loss from temptation. *Note:* The hatched region identifies types that would select into the public sector when $\lambda = 1$ but do not when $\lambda > 1$, because intermediate-γ agents incur self-control or commitment losses.

Proposition 3. *In the presence of corruption opportunities ($b > 0$), an increase in temptation λ:*

i) decreases the share of workers selecting into the public sector;

ii) increases the share of corrupt workers among public-sector employees;

iii) has an ambiguous effect on average ability and intrinsic motivation among public-sector entrants.

The first two statements follow from the shrinkage of the participation set in Figure 7 and from the expansion of the giving-in region as λ increases. The third statement is again compositional and depends on the joint distribution of (θ, γ).

Corollary 1. *If temptation is sufficiently strong (high λ), then the attractiveness of public employment is reduced for intrinsically motivated workers in the intermediate range, crowding them out of the public sector.*

Discussion. Temptation generates a wedge between corruption opportunities (which can attract low-γ types) and self-control costs (which repel intermediate-γ types). This mechanism implies that environments with stronger temptation may exhibit both a smaller public sector and a less honest workforce, even holding corruption opportunities fixed.

4. Benchmark analysis with uniform distributions

A central question in the literature on corruption and public employment is whether corruption opportunities primarily attract individuals with low intrinsic motivation or instead deter highly motivated and able workers. As discussed in Section 1, empirical evidence is mixed: in some contexts, pro-social preferences and ability reinforce each other (Barfort et al., 2019), while in others dishonesty and low motivation correlate with higher ability (Hanna and Wang, 2017). These contrasting findings mirror the ambiguity of the comparative statics derived in Section 3.

To discipline these predictions and clarify the economic forces at work, this section analyzes a set of benchmark cases under uniform distributions. The benchmark allows us to obtain closed-form expressions and to transparently link corruption opportunities and temptation to both the size and the composition of the public sector.

Assumptions. Throughout this section:

i) ability is uniformly distributed, $\theta \sim U[0, 1]$;

ii) public-service motivation $\gamma \in [0, 1]$ may be independent of θ, perfectly positively correlated, or perfectly negatively correlated;

iii) we normalize $w = 0$ and impose $\lambda b \leq 1$, so that all relevant cutoffs are interior.

We consider three benchmark correlation structures:

i) **Independence:** $(\gamma, \theta) \sim U[0, 1]^2$;

ii) **Perfect positive correlation:** $\gamma = \theta$;

iii) **Perfect negative correlation:** $\gamma = 1 - \theta$.

We begin with the independence case, which serves as a natural baseline and delivers closed-form results. We then turn to the two polar correlation structures to illustrate how alignment or misalignment between ability and motivation sharpens or weakens the models predictions.

4.1. Independence

Assume $(\gamma, \theta) \sim U[0,1]^2$ independently. Recall from Lemma 2 that the effective public-sector utility is

$$U_P(\gamma) = \begin{cases} b, & \gamma < \frac{1+\lambda}{2}b, \\ 2\gamma - \lambda b, & \gamma \in \left(\frac{1+\lambda}{2}b, \lambda b\right], \\ \gamma, & \gamma > \lambda b, \end{cases}$$

where $w = 0$.

Since private-sector utility equals θ and is independent of γ, entry into the public sector occurs whenever $\theta \leq U_P(\gamma)$.

Lemma 3. Selection probability: *With $(\gamma, \theta) \sim U[0,1]^2$, the probability that an individual with motivation γ enters the public sector equals $U_P(\gamma)$. Aggregate public-sector employment is therefore*

$$S_P = \int_0^1 U_P(\gamma)\, d\gamma.$$

This representation allows us to characterize the size of the public sector.

Proposition 4. Public-sector size: *Under the uniform and independent benchmark,*

$$S_P = \frac{1}{2} + \frac{1 + 2\lambda - \lambda^2}{4} b^2.$$

Public-sector employment increases in b if $\lambda < \lambda^\star$ and decreases in b if $\lambda > \lambda^\star$, where

$$\lambda^\star = 1 + \sqrt{2}.$$

The threshold $\lambda^\star$ coincides with the one identified in Corollary 1 of Section 3. In both cases, it reflects the balance between the attraction of corruption rents and the repulsion generated by temptation and self-control costs.

Changes in sector size are mirrored by changes in composition.

Proposition 5. Composition: *Under the uniform and independent benchmark, conditional on entry,*

$$\mathbb{E}[\gamma \mid P] = \frac{2}{3} + \frac{\lambda^2 - 2\lambda - 1}{3} b^2, \qquad \mathbb{E}[\theta \mid P] = \frac{1}{3} + \frac{\lambda^2 - 2\lambda - 1}{6} b^2.$$

Hence, both average motivation and ability decrease in b if $\lambda < \lambda^\star$ and increase in b if $\lambda > \lambda^\star$.

Finally, we can characterize the prevalence of corruption within the public sector.

Corollary 2. Corruption share: *The fraction of corrupt workers among public employees is*

$$f_C(b,\lambda) = \frac{\frac{1+\lambda}{2}b^2}{\frac{1}{2} + \frac{1+2\lambda-\lambda^2}{4}b^2},$$

which is strictly increasing in both b and λ.

Taken together, Proposition 4, Proposition 5 and Corollary 2 show that corruption opportunities have opposite effects depending on the strength of temptation. When temptation is mild ($\lambda < \lambda^\star$), the public sector expands but deteriorates in quality. When temptation is strong ($\lambda > \lambda^\star$), it contracts while becoming positively selected.

Remark 1. Positive wages: *The normalization $w = 0$ is adopted for clarity. Introducing a positive wage shifts all public-sector utilities upward without altering the qualitative comparative statics, as long as the relevant cutoffs remain interior.*

4.2. Perfect positive correlation

Assume $\gamma = \theta \sim U[0,1]$. In this case, selection depends on a single index, since ability and motivation move together. Entry into the public sector occurs whenever $\theta \leq U_P(\theta)$.

Proposition 6. *With perfect positive correlation, changes in corruption opportunities affect ability and motivation in the same direction. There exists a cutoff $\lambda^\star_+$ such that:*

i) if $\lambda < \lambda^\star_+$, increases in b expand the public sector but lower both average ability and motivation;

ii) if $\lambda > \lambda^\star_+$, increases in b shrink the sector while improving both margins.

Positive correlation therefore delivers sharp and unambiguous predictions: sector size and composition move one-for-one.

4.3. Perfect negative correlation

Finally, assume $\gamma = 1-\theta$ with $\theta \sim U[0,1]$. Entry requires $\theta \leq U_P(1-\theta)$. Any change in incentives that improves ability among entrants necessarily lowers motivation, and vice versa.

Proposition 7. *With perfect negative correlation, corruption opportunities generate an intrinsic trade-off between ability and motivation. While a threshold governs sector size as in the independent case, the effect on overall quality is fundamentally ambiguous.*

Remark 2. Correlation and ambiguity: *Positive correlation yields sharp predictions, while negative correlation destroys them. This helps reconcile why empirical studies often report mixed evidence on how rents affect the composition of the public workforce.*

Overall, the benchmark analysis shows that temptation mediates the impact of corruption opportunities on public-sector selection. Whether rents expand or contract the public sectorand whether they worsen or improve its compositiondepends jointly on the strength of temptation and on how ability and motivation are distributed in the population.

Why heterogeneity matters for policy. The correlation between ability and intrinsic motivation is not a modeling detail, but a structural feature of societies and labor markets that policymakers typically take as given. Educational systems, social norms, and early-career selection mechanisms jointly shape whether competence and integrity tend to reinforce each other or trade off. In societies where ability and honesty are negatively correlated, institutional reforms face an inherent constraint: policies that attract more skilled workers may systematically deter integrity, and vice versa. By contrast, when the correlation is positive, standard policy instruments—such as wage increases or monitoring—are more likely to improve both margins simultaneously.

This perspective helps explain why identical anti-corruption or wage policies can generate sharply different outcomes across institutional environments. Heterogeneity in the joint distribution of ability and motivation therefore defines the *feasible set* of policy outcomes, rather than merely affecting their magnitude.

5. General selection with temptation and self-control

The benchmark analysis in Section 4 delivers sharp intuition by exploiting closed-form expressions under uniform distributions. This section abstracts from those functional-form and distributional assumptions and characterizes occupational selection under temptation and self-control costs for arbitrary joint distributions of ability and intrinsic motivation.

We show that the core mechanisms identified in the benchmark are not an artifact of the uniform environment. Under weak regularity conditions, selection continues to be governed by the interaction between temptation, self-control costs, and the correlation between ability and honesty. In particular, we establish the existence of a critical institutional threshold that governs the sign of the compositional response to corruption opportunities, thereby clarifying the scope and robustness of the benchmark results.

5.1. Setup

Let (γ, θ) denote intrinsic motivation (honesty) and ability, respectively. An agent enters the public sector if and only if

$$\theta \leq s(\gamma; \lambda, w),$$

where $s(\gamma; \lambda, w)$ is the effective selection frontier induced by the public-sector utility defined in Section 2. We assume that $s(\cdot)$ is weakly increasing in γ, increasing in the public wage w, and nonincreasing in the temptation parameter λ. Let (γ, θ) have joint distribution F with bounded support and positive density.

Public-sector size is given by

$$M(\lambda, w) = \iint_{\theta \leq s(\gamma;\lambda,w)} f(\gamma, \theta)\, d\theta\, d\gamma,$$

and average intrinsic motivation among public employees is

$$\bar{\gamma}(\lambda, w) = \frac{\iint_{\theta \leq s(\gamma;\lambda,w)} \gamma f(\gamma, \theta)\, d\theta\, d\gamma}{M(\lambda, w)}.$$

5.2. A general tipping point

The next theorem characterizes how temptation affects selection independently of any parametric assumption on the distribution of types.

Theorem 1. General selection with temptation: *Suppose that: (i) the joint distribution of (γ, θ) has bounded support and positive density; (ii) the selection frontier $s(\gamma; \lambda, w)$ is continuous, weakly increasing in γ, increasing in w, and strictly decreasing in λ; (iii) increases in λ disproportionately reduce the attractiveness of public-sector employment for higher values of γ.*

Then:

i) Public-sector size $M(\lambda, w)$ is weakly decreasing in λ.

ii) There exists at least one threshold $\lambda^\star$ such that:

- *for $\lambda < \lambda^\star$, an increase in λ reduces average intrinsic motivation among public employees;*
- *for $\lambda > \lambda^\star$, an increase in λ raises average intrinsic motivation among public employees.*

iii) If the joint distribution satisfies a standard monotonicity condition (e.g. MLRP in γ), the threshold $\lambda^\star$ is unique.

Interpretation. The theorem establishes that temptation generically induces a non-monotonic compositional response. At low levels of temptation, increases in λ primarily attract marginal entrants or discourage highly motivated workers from entering, reducing average honesty. At sufficiently high levels of temptation, however, the least motivated types have already self-selected out of public employment, so further increases in λ act as a selective filter that improves composition. The threshold $\lambda^\star$ marks the point at which the marginal effect of temptation on selection changes sign.

Intuition. The result is driven by a geometric selection mechanism. Public-sector entry is governed by a frontier in the (γ, θ) space that shifts with institutional parameters. An increase in temptation does not affect all agents symmetrically: because resisting temptation is more costly for intrinsically motivated individuals, higher values of λ disproportionately reduce the attractiveness of public employment for high-γ types. At low levels of temptation, the rent-attraction effect dominates and corruption expands the public sector by drawing in marginal entrants. At sufficiently high levels of temptation, however, the selection margin flips: the agents who exit first are those with high intrinsic motivation, so further increases in λ reduce participation while improving average honesty among those who remain.

Proof sketch. The proof relies on selection through a moving frontier. Public-sector size and composition are obtained by integrating the joint density over the region below the selection frontier. An increase in λ shifts this frontier inward, with a larger displacement at higher values of γ under the maintained assumptions. As a consequence, the marginal mass leaving the public sector is initially concentrated among high-motivation agents, lowering average honesty. As λ increases further, the frontier enters low-density regions, attenuating this effect and eventually reversing its sign. Continuity of the frontier and bounded support guarantee the existence of a threshold $\lambda^\star$, while monotonicity conditions ensure uniqueness.

Formal arguments based on monotone-set inclusion and moving selection frontiers are developed in Appendix B, where we clarify the scope and robustness of the selection mechanism under arbitrary distributions.

Relation to the benchmark. Under the uniform and independent benchmark studied in Section 4, the threshold admits the closed-form expression $\lambda^\star = 1 + \sqrt{2}$. The general theorem shows that this value is not special: while the numerical location of the threshold depends on the distribution of types, its existence and qualitative role are robust.

6. Endogenizing temptation through institutional design

Thus far, the temptation parameter λ has been treated as an exogenous feature of the institutional environment. In practice, however, policymakers do not choose λ directly. Instead, they select observable institutional design levers—such as auditing intensity, monitoring, digitalization, or job rotation—that jointly affect both the scope for rent extraction and the burden of public employment.

This section provides a minimal and policy-relevant endogenization of temptation. Institutions choose a scalar design lever $m \geq 0$ that simultaneously: (i) reduces the effective temptation to engage in corruption, and (ii) raises organizational and psychological costs associated with public-sector work, such as compliance effort, reduced discretion, or procedural rigidity. The goal is not to solve an optimal policy problem, but to show that even this minimal mapping generates a sharp and empirically relevant trade-off.

6.1. Institutional mapping

We model institutional design as affecting temptation and self-control costs according to

$$\lambda(m) = 1 + \bar{\lambda}e^{-km}, \qquad \lambda'(m) < 0, \tag{8}$$

and

$$\psi(m) = \bar{\psi} + \eta m, \qquad \psi'(m) > 0, \tag{9}$$

where $\bar{\lambda} > 0$ captures baseline temptation in the absence of institutional effort, $k > 0$ governs how rapidly monitoring reduces temptation, and $\eta > 0$ measures the organizational burden induced by institutional design (e.g. compliance costs, reduced discretion, procedural rigidity).

The mapping reflects a realistic feature of anti-corruption reforms: policies that limit discretion and opportunities for rent extraction often increase the perceived cost of public employment, particularly for intrinsically motivated agents who would otherwise resist temptation.

6.2. Public-sector utility with institutional burden

Operationally, we incorporate institutional burden by modifying the effective public-sector utility in the self-control region. In the baseline model, agents who resist temptation incur a utility loss proportional to the temptation wedge. With institutional design, they also bear an additional burden $\psi(m)$.

The resulting effective utility is

$$U_m(\gamma) = \begin{cases} w + b, & \gamma < \frac{1+\lambda(m)}{2} b, \\ w + 2\gamma - \lambda(m)b - \psi(m), & \gamma \in \left(\frac{1+\lambda(m)}{2} b, \lambda(m)b\right], \\ w + \gamma, & \gamma > \lambda(m)b. \end{cases} \tag{10}$$

This formulation preserves the key mechanism of the baseline model. Increasing m lowers the temptation wedge (discipline channel, $\lambda(m) \downarrow$), but raises the cost of resisting temptation (burden channel, $\psi(m) \uparrow$).

6.3. Deterrence versus deterrence of entry

Let $S_P(m)$ denote the size of the public sector and $f_C(m)$ the fraction of corrupt workers among public employees. Both objects are defined as in Section 4, but now depend on m through $\lambda(m)$ and $\psi(m)$.

The next proposition summarizes the comparative statics.

Proposition 8. Institutional design and selection: *Assume* $(\gamma, \theta) \sim U[0,1]^2$ *i.i.d., normalize* $w = 0$*, and restrict to* $\lambda(m)b \leq 1$*. Then:*

i) The fraction of corrupt workers $f_C(m)$ *is weakly decreasing in* m *whenever* $\lambda'(m) < 0$*.*

ii) The effect of m *on public-sector size* $S_P(m)$ *is generally ambiguous: it increases through the discipline channel (*$\lambda(m) \downarrow$*) but decreases through the burden channel (*$\psi(m) \uparrow$*).*

iii) There exists a non-trivial region of parameters in which increasing m *reduces observed corruption while simultaneously reducing the entry of high-motivation agents.*

Proposition 8 highlights a fundamental institutional trade-off. Anti-corruption reforms that successfully deter corrupt behavior may also discourage participation by intrinsically motivated agents, shrinking the public sector and potentially worsening its composition.

Interpretation. This mechanism generates a sharp and testable prediction: environments with stronger monitoring and lower observed corruption may nevertheless exhibit lower public-sector participation or reduced attraction of highly motivated workers. The relationship linking monitoring intensity to occupational selection, together with the induced participation cutoff, as characterized in (8) and (9), provides a natural bridge between institutional design and empirically observable policy variables.

An alternative policy lever operates on the individual side rather than on enforcement: for example, cognitive behavioral training (CBT) or ethics programs can increase agents self-control capacity, effectively lowering the burden ψ for those in the temptation range, without altering monitoring intensity or sanctions.

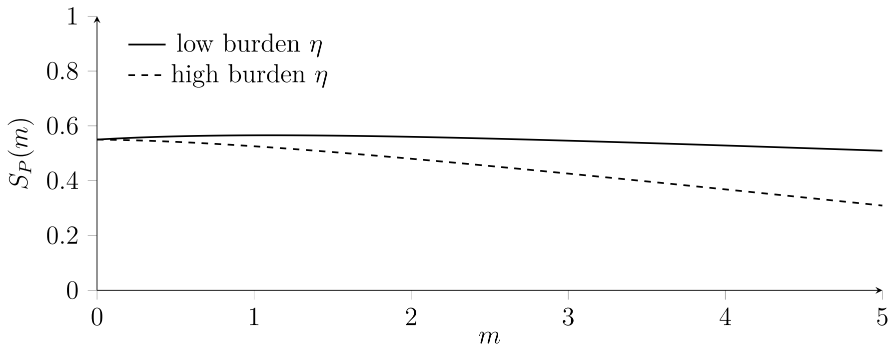

Fig. 8: Public-sector size as a function of institutional design. *Note:* m reduces temptation (discipline channel) but increases institutional burden (cost channel). When the burden slope η is large, $S_P(m)$ can decline even though corruption is deterred.

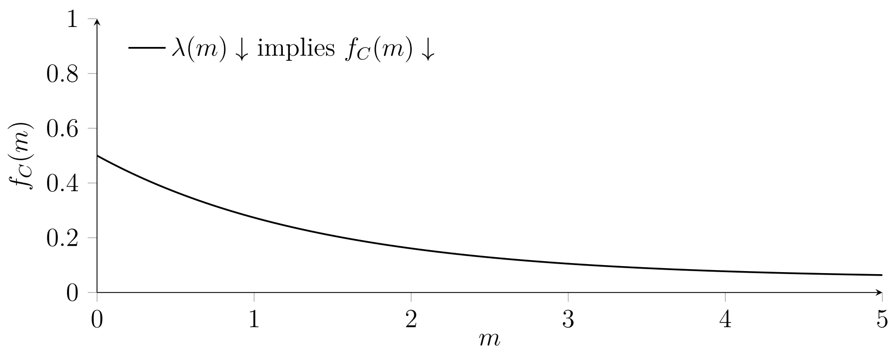

Fig. 9: Corruption share as a function of institutional design. *Note:* Increasing m reduces the effective temptation parameter $\lambda(m)$, shrinking the give-in region and lowering the fraction of corrupt among public employees.

7. Conclusion

This paper has developed a framework for analyzing occupational selection in environments characterized by moral frictions and self-control costs, taking corruption as the primary institutional friction. By integrating temptation and self-control into an otherwise standard selection model, the paper provides a unified explanation for the mixed empirical evidence on how corruption affects both the size and the quality of the public sector. Selection outcomes emerge from the interaction between behavioral costs, the joint distribution of ability and honesty in society, and the intensity of corruption opportunities.

The analysis delivers three main insights with implications for institutional design. First, in a tractable benchmark environment, we identify the existence of an *institutional tipping point* $\lambda^\star$ at which corruption opportunities switch from expanding public-sector participation to selectively deterring entry. This tipping point reflects a fundamental trade-off between the rent-attraction effect of corruption and the deterrence induced by self-control costs. Second, the results highlight the central role of the societal correlation between ability and honesty in shaping selection outcomes. When ability and honesty are negatively correlated, policymakers face an unavoidable trade-off: policies that improve one dimension of workforce quality may worsen the other. Third, the paper provides a novel interpretation of high public-sector wages. Beyond attracting skills, higher wages attenuate self-control costs and thereby shield motivated agents from the adverse selection effects of corrupt environments.

The general analysis developed in Section 5 clarifies the scope of these findings beyond the benchmark environment. While the numerical value of the tipping point $\lambda^\star$ is benchmark-specific, the underlying selection mechanism is geometric and persists under arbitrary distributions, correlation structures, and wage normalizations. Formal arguments based on monotone-set inclusion and moving selection frontiers are developed in Appendix B. More generally, policies that uniformly expand the selection frontier increase public-sector participation, whereas policies that disproportionately raise the cost of resisting temptation can reduce entry even as observed corruption declines.

Several avenues for future research remain open. One natural direction is to further discipline and optimize the institutional design channel developed here by linking the temptation parameter to specific policy instruments—such as monitoring intensity, auditing regimes, or bureaucratic discretion—in a fully calibrated or empirically grounded setting.

A complementary line of work would embed the selection mechanism into a political environment with electoral competition, such as a probabilistic voting model. This would allow the study of how corruption-induced selection interacts with policy platforms, voter heterogeneity, and electoral incentives, without altering the core behavioral mechanism analyzed in this paper. More broadly, extending the framework to a dynamic setting could shed light on how institutional environments interact with the formation and transmission of moral values over time.

By highlighting self-control costs as a key driver of labor-market sorting, this paper provides a theoretical foundation for interpreting and designing empirical studies on public-sector performance, corruption, and political selection.

References

[1] Acemoglu, D., Egorov, G., and Sonin, K. (2013). A Political Theory of Populism. *The Quarterly Journal of Economics*, 128(2):771–805.

[2] Acemoglu, D. and Verdier, T. (2000). The Choice between Market Failures and Corruption. *American Economic Review*, 90(1):194211.

[3] Banuri, S. and Keefer, P. (2016). Pro-Social Motivation, Effort and the Call to Public Service. *European Economic Review*, 83:139–164.

[4] Barfort, S., Harmon, N. A., Hjorth, F., and Olsen, A. L. (2019). Sustaining Honesty in Public Service: The Role of Selection. *American Economic Journal: Economic Policy*, 11(4):96–123.

[5] Barigozzi, F., Burani, N., and Raggi, D. (2018). Productivity Crowding-Out in Labor Markets With Motivated Workers. *Journal of Economic Behavior and Organization*, 151:199–218.

[6] Becker, G. S. (1968). Crime and Punishment: An Economic Approach. *Journal of Political Economy*, 76(2):169–217.

[7] Bénabou, R. and Tirole, J. (2006). Incentives and prosocial behavior. *American Economic Review*, 96(5):1652–1678.

[8] Bernheim, B. D. and Kartik, N. (2014). Candidates, Character, and Corruption. *American Economic Journal: Microeconomics*, 6(2):205–246.

[9] Besley, T. (2004). Joseph Schumpeter Lecture: Paying Politicians: Theory and Evidence. *Journal of the European Economic Association*, 2(2/3):193–215.

[10] Brassiolo, P., Estrada, R., Fajardo, G., and Vargas, J. (2021). Self-Selection into Corruption: Evidence from the Lab. *Journal of Economic Behavior and Organization*, 192:799–812.

[11] Bénabou, R. and Tirole, J. (2011). Identity, Morals, and Taboos: Beliefs as Assets. *The Quarterly Journal of Economics*, 126(2):805–855.

[12] Caselli, F. and Morelli, M. (2004). Bad Politicians. *Journal of Public Economics*, 88(3):759–782.

[13] Cervellati, M. and Vanin, P. (2013). "Thou Shalt not Covet: Prohibitions, Temptation and Moral Values. *Journal of Public Economics*, 103:15–280.

[14] Corneo, G. and Jeanne, O. (2009). A Theory of Tolerance. *Journal of Public Economics*, 93(5):691–702.

[15] Corneo, G. and Jeanne, O. (2010). Symbolic Values, Occupational Choice, and Economic Development. *European Economic Review*, 54:237–251.

[16] Dal Bó, E., Dal Bó, P., and Di Tella, R. (2006). “Plata o Plomo?": Bribe and Punishment in a Theory of Political Influence. *American Political Science Review*, 100(1):41–53.

[17] Dal Bó, E. and Di Tella, R. (2003). Capture by Threat. *Journal of Political Economy*, 111(5):11123–1154.

[18] Dal Bó, E., Finan, F., and Rossi, M. A. (2013). Strengthening State Capabilities: the Role of Financial Incentives in the Call to Public Service. *The Quarterly Journal of Economics*, 128(3):1169–1218.

[19] Delfgaauw, J. and Dur, R. (2010). Managerial Talent, Motivation, and Self-Selection into Public Management. *Journal of Public Economics*, 94(9):654–660.

[20] Ehrlich, I. (1973). Participation in Illegitimate Activities: A Theoretical and Empirical Investigation. *Journal of Political Economy*, 81(3):521–565.

[21] Gul, F. and Pesendorfer, W. (2001). Temptation and Self-Control. *Econometrica*, 69(6):1403–1435.

[22] Hanna, R. and Wang, S.-Y. (2017). Dishonesty and Selection into Public Service: Evidence from India. *American Economic Journal: Economic Policy*, 9(3):262–290.

[23] Konrad, K. A., Lohse, T., and Simon, S. A. (2021). Pecunia non Olet: on the SelfSelection into (Dis)Honest Earning Opportunities. *Experimental Economics*, 24:1105–1130.

[24] Lazear, E. P. (1979). Why Is There Mandatory Retirement? *Journal of Political Economy*, 87(6):1261–1284.

[25] Roy, A. D. (1951). Some Thoughts on the Distribution of Earnings. *Oxford Economic Papers*, 3(2):135–146.

Appendix

A. Proofs

Proof of Lemma 1 (Optimal behavior)

Proof. An agent employed in the public sector chooses between honesty (h) and corruption (c) by solving

$$\max_{x\in\{c,h\}} \left\{ u_{\mathrm{P}}(x) - \Big[\max_{y\in\{c,h\}} v_{\mathrm{P}}(y) - v_{\mathrm{P}}(x)\Big] \right\}.$$

This problem is equivalent to

$$x^* = \arg\max_{x\in\{c,h\}} \left\{ u_{\mathrm{P}}(x) + v_{\mathrm{P}}(x) \right\}.$$

Observe that $u_{\mathrm{P}}(c) > u_{\mathrm{P}}(h)$ if and only if $b > \gamma$, while $v_{\mathrm{P}}(c) > v_{\mathrm{P}}(h)$ if and only if $\lambda b > \gamma$. Three cases arise.

Case 1. If $\gamma > b$ and $\gamma > \lambda b$, honesty strictly dominates under both rankings.
Case 2. If $\gamma < b$ and $\gamma < \lambda b$, corruption strictly dominates under both rankings.
Case 3. If $\gamma \in [b, \lambda b]$, the comparison reduces to

$$u_{\mathrm{P}}(c) + v_{\mathrm{P}}(c) > u_{\mathrm{P}}(h) + v_{\mathrm{P}}(h) \quad \Leftrightarrow \quad \gamma < \tfrac{1+\lambda}{2} b.$$

This establishes the characterization of optimal behavior. □

Proof of Lemma 2 (Indirect public-sector utility)

Proof. We compute the indirect public-sector utility $U(\gamma)$ case by case.

If $\gamma < b$, there is no temptation and corruption is optimal, yielding $U(\gamma) = w + b$.
If $\gamma \in [b, \frac{1+\lambda}{2} b)$, corruption remains optimal and temptation is irrelevant, so $U(\gamma) = w + b$.
If $\gamma \in (\frac{1+\lambda}{2} b, \lambda b]$, the agent behaves honestly but incurs a self-control cost. Thus

$$U(\gamma) = u_{\mathrm{P}}(h) - \Big[v_{\mathrm{P}}(c) - v_{\mathrm{P}}(h)\Big] = w + 2\gamma - \lambda b.$$

If $\gamma > \lambda b$, there is again no temptation and honesty yields $U(\gamma) = w + \gamma$. □

Proof of Corollary 1 (Crowding-out threshold)

Proof. In the (γ, θ) plane, corruption opportunities generate a positive entry region of area $\frac{b^2}{2}$ and a negative exit region of area

$$\frac{(\lambda b - b)\left(\frac{1+\lambda}{2}b - b\right)}{2}.$$

Public-sector participation decreases if and only if the latter exceeds the former, i.e.

$$(\lambda b - b)\left(\tfrac{1+\lambda}{2}b - b\right) > b^2.$$

This inequality simplifies to $\lambda^2 - 2\lambda - 1 > 0$. Discarding the negative root yields the threshold $\lambda > 1 + \sqrt{2}$. □

Proof of Lemma 3 (Selection probability)

Proof. Since $\theta \sim U[0,1]$, the probability that $\theta \leq U(\gamma)$ equals $U(\gamma)$. Integrating over $\gamma \in [0,1]$ yields

$$S_P = \int_0^1 U(\gamma)\, d\gamma.$$

□

Proof of Proposition 5 (Composition under independence)

Proof. Expected public-service motivation conditional on entry is

$$\mathbb{E}[\gamma \mid P] = \frac{\int_0^1 \gamma U(\gamma)\, d\gamma}{S_P}.$$

Piecewise integration yields

$$\int_0^1 \gamma U(\gamma)\, d\gamma = \tfrac{1}{3} + \left(-\tfrac{\lambda^3}{8} + \tfrac{\lambda^2}{8} + \tfrac{\lambda}{8} + \tfrac{1}{24}\right)b^3,$$

so that

$$\mathbb{E}[\gamma \mid P] = \tfrac{2}{3} + \tfrac{\lambda^2 - 2\lambda - 1}{3}b^2 + O(b^3).$$

Expected ability conditional on entry satisfies $\mathbb{E}[\theta \mid \theta \leq U(\gamma)] = U(\gamma)/2$. Hence

$$\mathbb{E}[\theta \mid P] = \frac{\int_0^1 \frac{1}{2}U(\gamma)^2\, d\gamma}{S_P} = \tfrac{1}{3} + \tfrac{\lambda^2 - 2\lambda - 1}{6}b^2 + O(b^3).$$

The sign of both effects is governed by $\lambda^2 - 2\lambda - 1$. □

Proof of Corollary 2 (Corruption share)

Proof. The mass of corrupt entrants equals

$$M_C = \int_0^a b\, d\gamma = \tfrac{1+\lambda}{2} b^2, \quad \text{where } a = \tfrac{1+\lambda}{2} b.$$

Dividing by public-sector size yields

$$f_C(b, \lambda) = \frac{\tfrac{1+\lambda}{2} b^2}{\tfrac{1}{2} + \tfrac{1+2\lambda-\lambda^2}{4} b^2}.$$

Straightforward differentiation shows that f_C is increasing in both b and λ. $\square$

Proof of Proposition 6 (Positive correlation)

Proof. If $\gamma = \theta \sim U[0, 1]$, public-sector size is

$$S_P = \int_0^1 U(x)\, dx, \quad x = \gamma = \theta.$$

Partition $[0, 1]$ into $[0, a]$, $[a, c]$, and $[c, 1]$, with $a = \frac{1+\lambda}{2} b$ and $c = \lambda b$. Then

$$S_P = \int_0^a b\, dx + \int_a^c (2x - \lambda b)\, dx + \int_c^1 x\, dx = \tfrac{1}{2} + \tfrac{1+2\lambda-\lambda^2}{4} b^2.$$

Since $\gamma = \theta$, average motivation and ability coincide. Below $\lambda^\star = 1 + \sqrt{2}$ the sector expands while quality deteriorates; above $\lambda^\star$ it contracts while quality improves. $\square$

Proof of Proposition 7 (Negative correlation)

Proof. If $\gamma = 1 - \theta$ with $\theta \sim U[0, 1]$, marginally $\gamma \sim U[0, 1]$ and

$$S_P = \int_0^1 U(1 - \theta)\, d\theta = \tfrac{1}{2} + \tfrac{1+2\lambda-\lambda^2}{4} b^2.$$

Composition is given by

$$\mathbb{E}[\theta \mid P] = \frac{\int_0^1 \theta U(1-\theta)\, d\theta}{S_P}, \qquad \mathbb{E}[\gamma \mid P] = \frac{\int_0^1 (1-\theta) U(1-\theta)\, d\theta}{S_P}.$$

By symmetry, these expectations move in opposite directions with b, so one increases while the other decreases, yielding an ambiguous effect on overall quality. $\square$

Proof of Proposition 8 (Institutional design)

Proof. Under $(\gamma, \theta) \sim U[0,1]^2$ i.i.d. and $w = 0$, entry conditional on γ satisfies $\Pr(\theta \leq U_m(\gamma)) = U_m(\gamma)$, so $S_P(m) = \int_0^1 U_m(\gamma)\, d\gamma$.

(i) Corruption share. An agent behaves corruptly if and only if $\gamma < \frac{1+\lambda(m)}{2} b$. Thus the mass of corrupt entrants is

$$M_C(m) = \int_0^{a(m)} b\, d\gamma = a(m)b, \quad a(m) = \tfrac{1+\lambda(m)}{2} b.$$

Since $\lambda'(m) < 0$, we have $a'(m) < 0$ and $M_C(m)$ decreases in m. For interior solutions, this implies that the corruption share $f_C(m) = M_C(m)/S_P(m)$ is weakly decreasing.

(ii) Sector size. Changes in m affect $S_P(m)$ through two channels: (i) a discipline channel, as $\lambda(m)$ shifts the boundaries of the give-in and self-control regions; (ii) a burden channel, as $\psi(m)$ lowers utility in the self-control region. The net effect on $S_P(m)$ is therefore ambiguous.

(iii) Existence of a trade-off. For parameters such that $\psi'(m)$ is sufficiently large relative to $|\lambda'(m)|$, the burden channel dominates and $S_P(m)$ falls, while $M_C(m)$ continues to fall. In this region, increasing m reduces observed corruption but also reduces entry and can worsen selection. $\square$

B. Extensions and Robustness

This appendix collects robustness results that clarify which forces in the model are *purely geometric*arising from selection through a monotone entry frontier and which rely on the benchmark assumptions adopted in Section 4. The purpose of this appendix is not to extend the model, but to delineate the scope and generality of the selection mechanisms highlighted in the main text.

B.1. Setup and notation

Let γ denote intrinsic motivation (honesty) and θ ability. An agent enters the public sector if and only if

$$\theta \leq s(\gamma; z),$$

where $s(\cdot; z)$ is weakly increasing in γ and depends on an institutional vector z (e.g. public wages, corruption opportunities, or temptation-related parameters). Let (γ, θ) have joint cumulative distribution function F with bounded support $\mathcal{S}$ and density f.

Define the public-sector region

$$\mathcal{R}(z) = \{(\gamma, \theta) \in \mathcal{S} : \ \theta \leq s(\gamma; z)\},$$

sector size

$$M(z) = \iint_{\mathcal{R}(z)} f(\gamma, \theta)\, d\theta\, d\gamma,$$

and average intrinsic motivation among entrants

$$\bar{\gamma}(z) = \frac{\iint_{\mathcal{R}(z)} \gamma f(\gamma, \theta)\, d\theta\, d\gamma}{M(z)}.$$

Remark (Wage normalization). Section 4 normalizes the public wage to $w = 0$ for tractability. Allowing for $w > 0$ shifts the entry frontier upward without altering the selection logic, although it may attenuate comparative statics by expanding participation into regions where most types enter.

B.2. Robust comparative statics from monotone-set inclusion

Lemma B.1. Monotone-set inclusion: *Fix a distribution F with positive density on bounded support. If $s(\gamma; z') \geq s(\gamma; z)$ for all γ, then $\mathcal{R}(z) \subseteq \mathcal{R}(z')$ and therefore $M(z') \geq M(z)$.*

Proof. If $s(\gamma; z') \geq s(\gamma; z)$ for all γ, then any type satisfying $\theta \leq s(\gamma; z)$ also satisfies $\theta \leq s(\gamma; z')$. Hence $\mathcal{R}(z) \subseteq \mathcal{R}(z')$. Integrating a nonnegative density over a larger set yields $M(z') \geq M(z)$. $\square$

Implication. Any policy instrument that shifts the entry frontier uniformly outward (e.g. higher public wages) increases public-sector size independently of the specific shape of the joint distribution f. By contrast, instruments that affect the frontier non-uniformly in γ need not deliver monotone effects on participation or composition.

B.3. The uniform benchmark: the tipping point and its interpretation

Section 4 provides closed-form expressions under $(\gamma, \theta) \sim U[0, 1]^2$ and $w = 0$. In that benchmark,

$$S_P(b, \lambda) = \frac{1}{2} + \frac{1 + 2\lambda - \lambda^2}{4}\, b^2,$$

so that the effect of corruption opportunities b on public-sector size changes sign at

$$\lambda^{\star} = 1 + \sqrt{2}.$$

For $\lambda < \lambda^\star$, higher corruption opportunities expand the public sector; for $\lambda > \lambda^\star$, they contract it. This tipping point reflects the core trade-off of the model: corruption rents attract entry, while temptation-related self-control costs deter intrinsically motivated agents.

Remark (What is robust and what is benchmark-specific). The existence of a tipping point is driven by the piecewise structure of the effective public-sector utility and the geometry of the induced entry set. Outside the uniform benchmark, the numerical value of $\lambda^\star$ need not equal $1 + \sqrt{2}$. However, the underlying mechanisma competition between an attraction force (rents) and a deterrence force (self-control costs)is general.

B.4. Correlation and continuity

To study intermediate correlation patterns between ability and motivation, let the joint distribution be constructed via a copula:

$$F_\rho(\gamma, \theta) = C_\rho(F_\gamma(\gamma), F_\theta(\theta)), \qquad \rho \in (-1, 1),$$

with fixed marginals F_γ and F_θ and a parametric family of copulas C_ρ.

Proposition B.1. Continuity in the correlation parameter: *Assume that $s(\gamma; z)$ is continuous in (γ, z) and that the copula density exists and varies continuously in ρ. Then $M(z; \rho)$ and $\bar{\gamma}(z; \rho)$ are continuous in ρ.*

Proof. Continuity follows from dominated convergence: both M and the numerator defining $\bar{\gamma}$ are integrals of bounded functions over bounded support with densities that vary continuously in ρ. □

Takeaway. The polar cases analyzed in Section 4 are informative endpoints. Intermediate correlation structures generate intermediate outcomes smoothly, with no discontinuities as ρ varies.

B.5. High wages attenuate temptation effects

High public wages shift the entry frontier outward and reduce the sensitivity of participation to temptation-related wedges by bringing a broader set of types into the public sector.

Proposition B.2. Attenuation under large wages: *Suppose $s(\gamma; w, \lambda)$ is increasing in w and nonincreasing in λ for all γ. Then, for any fixed λ, the marginal impact of λ on participation weakly decreases as w increases, in the sense that the* *change in $M(w, \lambda)$ induced by a given increase in λ becomes smaller when w is larger.*

This attenuation result formalizes a simple intuition: when public employment is already attractive due to high pay, temptation-related deterrence operates on a smaller margin of agents.

B.6. Empirical dictionary

This dictionary is intended as a guide for empirical interpretation rather than a one-to-one identification strategy (see also Section 1 and Section 4).

Model primitive	Observable proxy
γ (intrinsic honesty)	Dictator-game giving, honesty tasks
θ (ability)	Exam scores, cognitive tests
b (corruption rents)	Bribe exposure, illicit gain size
λ (temptation intensity)	Monitoring intensity, payoff immediacy
m (institutional design)	Audit frequency, procedural rigidity

Table 1: Mapping between model primitives and standard experimental or administrative measures used in the empirical literature.

Crime Perception and Voting Behavior: Evidence from Individual Data*

Giovanni Prarolo[†] Marco Rosso[†]

Abstract

This paper studies how exposure to geolocated crime-related news shapes individual voting behavior in Italian elections. Using a panel of non-relocating voters observed across multiple election rounds, we exploit within-individual variation in exposure to nearby crime news during the pre-election month, controlling for individual and district-by-election fixed effects. Aggregate exposure yields weak, unstable effects. Disaggregating by offender nationality reveals systematic patterns: immigrant-attributed crime news reduces support for parties with ambiguous immigration stances (e.g., Five Star Movement) and increases support for clear "law-and-order" parties, while Italian-attributed crime has negligible effects. Effects are stronger among high-skilled voters shifting from M5S and low-skilled voters abandoning Lega. In local elections, Italian crime punishes incumbents, while immigrant crime increases abstention. These asymmetric responses—absent in aggregate measures—indicate that crime salience operates primarily through identity-based framing rather than through generalized concerns about crime or security. The findings highlight how media attribution shapes electoral accountability.

Keywords: crime; immigration; elections; news media; individual voting behavior.

JEL: D72; D83; K42; L82.

*We thank participants in seminars and conferences in Bertinoro, Bologna, Bruneck, Buenos Aires, Modena, Naples, and Thessaloniki for their helpful comments. We are especially grateful to Domenico Depalo, Giovanni Mastrobuoni, Nicola Mastrorocco, Tommy E. Murphy, Antonio Schiavone, Federico Trombetta, Orestis Troumpounis, Juan F. Vargas, and Paolo Vanin for their detailed feedback. We gratefully acknowledge financial support from Fondazione Carisbo and from the Italian Ministry of University (PRIN projects 2017KHR4MB and 2022SHHM2A). All errors are our own.

†Prarolo: Department of Economics, University of Bologna (email: giovanni.prarolo@unibo.it); Rosso: Department of Economics, University of Bologna (email: marco.rosso4@unibo.it).

1. Introduction

Across many European democracies, and especially in Italy, debates over immigration have increasingly been framed through the lens of public safety. A large body of evidence shows that citizens often associate immigration with crime even when such links are weak or distorted in the data, and that media reporting can amplify this association by selectively making certain events salient (Keita et al., 2023; Ajzenman et al., 2023; Couttenier et al., 2024). These dynamics are politically consequential: when crime is framed in ways that emphasize the identity of the perpetrator, it can strengthen narratives of "threat" and shift electoral support toward parties promising tougher enforcement or stricter immigration policies. More generally, research on persuasion highlights that changes in voters information sets and beliefs can translate into changes in political behavior (DellaVigna and Gentzkow, 2010; DellaVigna and Kaplan, 2007; Enikolopov et al., 2011; Gerber et al., 2009; Barone et al., 2015).

Yet an important question remains unresolved. Do voters primarily react to crime itself, or does crime framed as immigrant-related trigger distinct political responses? The distinction matters because it separates two channels that are often conflated in political discourse. One channel is general crime salience: learning that crime occurred nearby may increase concern about safety and raise support for "law and order" platforms. A second channel is identity framing: when crime is attributed to immigrants, it may also affect attitudes toward immigration and minorities, consistent with mechanisms rooted in stereotypes and availability-driven belief formation (Kahneman and Tversky, 1973; Bordalo et al., 2018). Disentangling these channels is essential for understanding how local shocks to perceived security translate into shifts in party support and the electoral incentives faced by political actors. Crucially, we argue that crime-related news affects voting behavior not through mechanical responses to local crime risk, but only when reporting activates identity-based political narratives—most notably those linking immigration to law and order.

This paper investigates whether exposure to salient information about immigrant-related crime in ones immediate neighborhood shifts individual voting behavior differently from exposure to comparable crimes attributed to Italian nationals. We study this question in Bologna, a large city in Northern Italy, where voting behavior can be measured at the individual level over repeated elections and where local crime information is disseminated through a prominent online outlet. Our key empirical object is *local crime news exposure*. Specifically, we define exposure as the publication of a crime-related news report that can be geocoded within 200 meters of an individuals home address during the month preceding an election round. Crimes attributed to Italian nationals serve as a benchmark that allows us to separate general crime salience from the immigration-specific framing channel emphasized by populist rhetoric (Keita et al., 2023; Ajzenman et al., 2023).

We interpret our measure as capturing the subset of criminal events that becomes *salient information* to residents through news reporting. We do not claim that media coverage is a complete or unbiased measure of the universe of crimes; rather, it captures the information set plausibly available to voters in real time and at street-level granularity, which is the relevant object for belief formation and political responses (DellaVigna and Gentzkow, 2010; Mastrorocco and Minale, 2018). Unlike much of the existing literature, our approach embeds fine-grained spatial and temporal variation directly by linking individuals to geocoded local crime reports, rather than relying on cross-sectional variation in media access combined with aggregate time variation in news content.[1]

Our empirical design compares declared voting choices across multiple elections for the same individuals who are differentially exposed to nearby immigrant-related and native-related crime reports. The analysis exploits fine spatio-temporal variation in the timing and location of local crime news, allowing us to track individual switching behavior across elections. Conceptually, the strategy follows a differences-in-differences design that leverages within-individual variation combined with plausibly idiosyncratic shocks in local crime reporting at the street level.

To implement this design, we combine two novel data sources. First, we use a representative survey of 5,000 Bologna residents conducted in 2021 that records retrospective voting behavior across nine elections between 2004 and 2021. Second, we assemble a corpus of approximately 11,000 crime-related articles published between 2011 and 2022 by *BolognaToday*. Using a dictionary-based coding scheme standard in text-as-data applications (Gentzkow et al., 2019), we classify whether reports describe perpetrators as immigrants or Italian nationals and geocode each report at street level. The underlying behavioral premise is straightforward: information about a crime occurring in ones immediate vicinity is more likely to be noticed and perceived as relevant than information about crime occurring farther away, consistent with salience and availability mechanisms (Kahneman and Tversky, 1973; Mastrorocco and Minale, 2018).

Because our empirical strategy relies on local media reporting, a natural concern is whether coverage itself responds to political preferences across space. We address this concern directly. Exploiting a difference-in-discontinuities design in the spirit of Grembi et al. (2016), we compare crime coverage across historically right- and left-leaning neighborhoods around electoral thresholds and find no evidence of

[1] This differs from studies exploiting the staggered introduction of Fox News in the United States (DellaVigna and Kaplan, 2007), differential access to independent television in Russia (Enikolopov et al., 2011), or the expansion of digital television channels in Italy (Barone et al., 2015).

differential coverage or framing of crime by area ideology.

Our main findings reveal asymmetric electoral effects. Individuals exposed to nearby immigrant-related crime reports become significantly more likely to shift their support toward right-wing and anti-immigration parties, both in national and local elections. By contrast, exposure to comparable crimes attributed to Italian nationals does not generate similar shifts. In national elections, immigrant-related crime exposure reduces switching toward the Center-Left and toward the Lega, while increasing switching toward the broader Center-Right coalition. In local elections, immigrant-related crime exposure reduces switching toward the Center-Left while increasing abstention and support for smaller parties, with additional evidence of mobilization toward the Center-Right.

This paper contributes to several strands of literature. First, it adds to research in the economics of crime documenting that crime has broader consequences beyond victimization itself, affecting economic and political outcomes (Buonanno et al., 2013; Dustmann and Fasani, 2016; Drago et al., 2020). Second, it contributes to the literature on media and political persuasion, which shows that exposure to media content can shift beliefs and electoral outcomes (DellaVigna and Kaplan, 2007; Enikolopov et al., 2011; Gerber et al., 2009; Barone et al., 2015). Our evidence further complements work documenting partisan distortions in Italian television news (Durante and Knight, 2012) and their electoral consequences (Barone et al., 2015). Related research shows that crime coverage in particular shapes perceptions of safety and attitudes toward immigrants (Mastrorocco and Minale, 2018; Ajzenman et al., 2023; Sacco, 1982; Couttenier et al., 2024).

Third, our results speak to a broader literature on stereotypes and misperceptions, which shows that vivid and salient signals can reinforce distorted beliefs about immigration (Alesina et al., 2023). Our findings are consistent with mechanisms in which immigrant-related crime news disproportionately affects beliefs through availability and stereotype-based distortions (Kahneman and Tversky, 1973; Bordalo et al., 2018). Finally, our analysis relates to work on crime, social dynamics, and political participation (Blanes i Vidal and Mastrobuoni, 2018; Kirchmaier et al., 2020; Vargas et al., 2025).

In sum, we provide novel evidence that exposure to salient information about immigrant-related crime in ones immediate environment has a significant impact on voting behavior. By linking individual voting histories to geocoded local crime reporting, we offer a micro-level perspective on how identity framing and crime salience interact in shaping electoral outcomes. The remainder of the paper proceeds as follows. Section 3 describes the data and the construction of exposure measures. Section 3 outlines the empirical strategy. Section 5 presents the main results and heterogeneity analyses, while Section 5 provides supporting evidence for the identification strategy. Section 6 concludes. Additional information and robustness checks are provided in the Appendix.

2. Data

2.1. Survey Data

Our analysis relies on the survey data originally designed and collected for the work Berti Ceroni et al. (2025), so most of the technical details can be found therein. We will stress here the fundamental features that are key to this analysis.

The survey was administered in December 2021 by a professional firmand covers 5,000 Italian citizens residing in Bologna, a 400-thousands-inhabitants city in Northern Italy, who were eligible to vote and had lived in the city since at least 2013. Respondents were stratified across the citys 18 neighborhoods, with interviews conducted via telephone (70%) and online (30%). The questionnaire collected detailed demographic and socio-economic characteristics—such as birth year and province, education, employment status, occupation, and family structure and, key to our purposes, retrospective voting choices in all national (2006, 2008, 2013, 2018) and municipal elections (2004, 2009, 2011, 2016, 2021).[2] A distinctive feature of the dataset obtained from the survey is the provision of respondents street of residence (without house numbers for privacy). This information enables sufficiently precise geo-localization and makes it possible to link individual voters to spatially disaggregated measures of media coverage of crime in their neighborhood.

To enhance recall accuracy regarding retrospective voting, a brief neutral summary of the political environment around each specific electoral round was provided to the participants before asking about their voting behaviors. Moreover, to mitigate potential consistency bias, voting questions were asked in chronological order, starting from the earliest election.

Validation against official electoral outcomes shows close alignment, with minor discrepancies only in the 2011 and 2016 municipal elections, when the Center-Right and Lega presented joint candidates. Internal mobility within Bologna is limited: more than 96% of respondents in 2021 lived in the same area as in 2013, reducing concerns about voting-with-feet within the city boundaries. Additional descriptive evidence and validation exercises are presented in Appendix A.

[2]The empirical analysis will exploit five electoral rounds: the national elections of 2013 and 2018, and the municipal elections of 2011, 2016, and 2021.

2.2. Crime Measures from Newspapers

Our primary source of crime data is *BolognaToday*, a leading independent online newspaper.[3] According to *Similarweb* data, the outlet attracts approximately 600 thousand monthly visitors from the Bologna metropolitan area, making it a highly relevant source of local information for the electorate.[4] To construct the dataset, we implemented an automated web scraping protocol targeting the newspapers archive section, which is dedicated to local news and daily events. We extracted the complete digital archive of this section from its establishment in 2011 through 2022. The scraper parsed the HTML structure of the website to systematically retrieve the headline, summary, full text, and publication timestamp for approximately 55,000 unique articles.

We classify the content of these articles using a dictionary-based algorithm designed to minimize false positives (Gentzkow et al., 2019; Muço, 2025). Rather than relying on simple keyword matching, our approach identifies crime events based on the co-occurrence of specific crime-category keywords (e.g., *theft*, *assault*, *drugs*, *murder*) with terms indicating a verified police action or report (e.g., "*arrested*", "*denounced*", "*investigated*"). Simultaneously, the algorithm scans for linguistic markers distinguishing the nationality of the offenders (Italians vs. non-Italians). In line with standard conventions in local crime reporting, references to non-Italian nationality are made explicit, whereas crimes committed by Italian citizens are typically reported without nationality markers; accordingly, articles that do not specify nationality are classified as involving Italian offenders. The whole procedure identifies crime-related news in 18.5% of the local news archive. To assess the reliability of the automated nationality classification, we manually coded a random sample of 100 crime-related articles; the resulting confusion matrix shows high accuracy, precision, and recall, indicating that misclassification is limited and unlikely to drive the main results (see Table B.1).

We successfully geolocated approximately 76% of these crime reports at the street level. It is important to note that serious offenses such as murders or kidnappings represent a marginal fraction of this coverage—roughly 3.7% of crime articles. Because high-profile violent crimes generate highly repeated coverage, the effective number of unique violent events is likely even lower; consequently, the dataset is dominated by less violent categories such as theft or assault.[5] The resulting daily

[3] At the time of data collection, the website was freely accessible. Subscription plans for premium content were introduced only in 2025, after the period considered in our analysis.

[4] To put this magnitude in perspective, this readership creates a ratio of roughly one visitor for every two residents in the metropolitan area of Bologna.

[5] According to the Italian National Institute of Statistics (ISTAT), the annual volume of minor property and violent crimes has averaged roughly 3 million in recent years, relative to 350 homicides.

panel dataset consists of approximately 7,800 geolocated crime articles, of which 46% involve immigrant offenders. To address potential measurement error arising from redundant reportingwhere a single incident generates multiple articles over several days—we do not use raw article counts as our primary measure of crime intensity. Instead, we construct indicator variables equal to one if a crime is reported in a specific location on a given day. These geocoded indicators form the basis for constructing individual-level measures of exposure to crime news, which we merge with electoral records. Further technical details on the dictionary validation and scraping architecture are provided in Appendix B.

3. Empirical Strategy

The empirical analysis examines whether exposure to *crime-related news* in the immediate vicinity of an individual affects voting behavior in national and local elections. We focus on crime-related news coverage referring to criminal events occurring within a 200-meter radius from the respondents street of residence during the 30 days preceding each election. This time window coincides with the official electoral campaign period regulated by the Italian "*par condicio*" law (Law 28/2000), which governs media coverage and ensures equal access to public and private media outlets in the run-up to elections.

Throughout this section, we use the term "crime-related news" to denote media coverage of criminal events, rather than underlying crime incidence.

Our contribution lies in distinguishing between crime-related news attributed to Italian citizens and those attributed to immigrants, allowing us to assess whether voters respond differently depending on the perceived nationality of offenders. The data are organized as an individual-level panel, where the cross-sectional dimension is the voter and the time dimension corresponds to successive election rounds.

To estimate the effect of exposure to crime-related news on voting behavior, we rely on fixed-effects specifications of the following form:

$$Vote_{it}^{P} = \alpha CI_{it} + \beta CNI_{it} + \gamma_i + \delta_t + \epsilon_{it}, \tag{1}$$

$$Vote_{it}^{P} = \alpha CI_{it} + \beta CNI_{it} + \gamma_i + \delta_{dt} + \epsilon_{it}, \tag{2}$$

where $Vote_{it}^{P}$ is a dummy equal to one if individual i votes for party (or option) P in election t, and zero otherwise. Elections are grouped by type (national and administrative), and we estimate separate regressions for each party or voting option, including abstention.

The variables CI_{it} and CNI_{it} are indicators equal to one if at least one crime-related news item attributed to Italian citizens or immigrants, respectively, is reported within 200 meters of the individuals residence during the 30 days preceding election t. We rely on binary exposure measures rather than article counts because a single criminal incident may generate multiple news reports over several days. Using indicator variables mitigates distortions arising from repeated coverage of the same event and focuses on the salience of crime-related news exposure rather than its volume.[6]

Accordingly, we interpret these indicators as measures of exposure to salient crime-related news rather than to underlying crime incidence.

The term γ_i denotes individual fixed effects, which absorb all time-invariant voter characteristics, including baseline political preferences, ideology, and residential location. The fixed effects δ_t capture election-specific shocks common to all voters, while δ_{dt} denotes district-by-election fixed effects that flexibly control for time-varying local political and socioeconomic conditions.

Identification of α and β relies on within-individual variation in exposure to crime-related news across elections, conditional on election-wide and district-specific shocks, and individual fixed effects. This variation arises from the idiosyncratic timing and highly localized spatial nature of crime-related news coverage within districts across election cycles. Intuitively, the coefficients capture how changes in local crime-news exposure relative to an individuals own past elections are associated with changes in the probability of voting for a given party. Because individual fixed effects are included, the estimates do not reflect cross-sectional differences in voting levels, but rather within-individual changes in vote choice relative to previous elections, holding constant baseline political preferences.

Standard errors are clustered at the district level to account for spatial correlation in crime reporting and local media coverage.

Two identification concerns merit discussion. First, crime-news coverage may be strategically manipulated around elections or may correlate with the historical political orientation of neighborhoods. Second, exposure to crime-related news may display short-run persistence around election dates that is not fully absorbed by the fixed effects structure. We address these concerns through a set of diagnostic exercises discussed in Section 5.

Specifically, we implement a difference-in-discontinuities analysis around national elections to test for discontinuous changes in crime-news coverage across districts with different historical political orientations (Subsection 5.1). In addition, we estimate placebo regressions using crime-related news occurring after election dates. As shown in Subsection 5.2, these exercises reveal no evidence of strategic report-

[6] Appendix C reports robustness checks using alternative radii (300 and 500 meters) and specifications based on article counts.

ing around elections or of spurious post-election effects. Together, they mitigate concerns related to strategic reporting and short-run persistence, supporting the interpretation of crime-related news exposure as plausibly exogenous with respect to individual vote switching within the election cycle.

4. Main Results

Our primary analysis focuses on national elections. Concentrating on national contests mitigates the role of local political issues and institutional idiosyncrasies, and allows us to study voting behavior in a setting where parties compete on national platforms and where law-and-order and immigration are salient dimensions of political debate.

Table 2 reports the main results. Tab. 1A presents estimates based on an aggregate measure of crime-related news exposure, while Tab. 1B distinguishes between crime-related news attributed to Italian citizens and to immigrants.

Table 1: Effects of Crime Perception on Voting at National Elections

Panel A: Overall Crimes

Dep.	M5S (1)	M5S (2)	Lega (1)	Lega (2)	Center Right (1)	Center Right (2)	Center Left (1)	Center Left (2)	Abstention (1)	Abstention (2)
Total Crimes	-0.012**	-0.011*	-0.005	-0.005	0.010	0.010	0.007**	0.004**	0.007**	0.007**
	(0.004)	(0.005)	(0.010)	(0.012)	(0.006)	(0.006)	(0.003)	(0.001)	(0.003)	(0.002)
Individual fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓		✓		✓		✓		✓	
Districts fixed effects × Year fixed effects		✓		✓		✓		✓		✓
R^2	0.05	0.05	0.16	0.16	0.11	0.11	0.02	0.02	0.01	0.01
Observations	9624	9624	9624	9624	9624	9624	9624	9624	9624	9624

Panel B: Crimes by Nationality of the Offenders

Dep.	M5S (1)	M5S (2)	Lega (1)	Lega (2)	Center Right (1)	Center Right (2)	Center Left (1)	Center Left (2)	Abstention (1)	Abstention (2)
Crimes by Italians	-0.004	-0.004	0.000	-0.002	0.002	0.004	0.004	0.002	0.005	0.007
	(0.007)	(0.007)	(0.014)	(0.014)	(0.013)	(0.013)	(0.003)	(0.003)	(0.005)	(0.005)
Crimes by immigrants	-0.023**	-0.021**	-0.015**	-0.010	0.021**	0.021**	0.007	0.008	0.005	0.001
	(0.007)	(0.007)	(0.004)	(0.007)	(0.006)	(0.008)	(0.006)	(0.006)	(0.010)	(0.011)
Individual fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓		✓		✓		✓		✓	
Districts fixed effects × Year fixed effects		✓		✓		✓		✓		✓
R^2	0.05	0.05	0.16	0.16	0.11	0.11	0.02	0.02	0.01	0.01
Observations	9624	9624	9624	9624	9624	9624	9624	9624	9624	9624

Notes: This table reports OLS estimates of the effect of local crime on voting behavior at national elections. The unit of observation is the individual i at time t. The dependent variable is a dummy that takes the value 1 if the individual votes for parties such as M5S, Lega, Center Left, Center Right, or abstains from voting. In Tab. 1A the independent variable is a dummy that takes value 1 if there is at least one crime within 200 meters of the respondents street in the 30 days preceding the election day. In Tab. 1B the key explanatory variables are indicators for whether at least one crime is committed by Italians or immigrants within 200 meters of the respondents street in the 30 days preceding the election day. Robust standard errors are clustered at district level.

Aggregate crime exposure. Tab. 1A shows that when crime-related news exposure is measured in the aggregate—without distinguishing the nationality of the offenders—the estimated effects on voting behavior are small and unstable across parties. While some coefficients are statistically significant, their magnitudes are modest and their signs vary across political options, offering no coherent pattern of electoral response. This lack of structure suggests that aggregate measures of crime-related news obscure important heterogeneity in how voters interpret and react to crime information.

Crime-related news by offenders' nationality. Tab. 1B reveals a markedly different and substantially more structured pattern once crime-related news is disaggregated by the perceived nationality of offenders. Exposure to crime-related news attributed to Italian citizens has no statistically significant effect on voting behavior across parties. In contrast, exposure to crime-related news attributed to immigrants induces systematic changes in individual vote choices.

Specifically, immigrant-related crime news is associated with a decline in the probability of voting for the Five Star Movement (M5S), with estimated effects ranging between 2.1 and 2.3 percentage points. We also observe a smaller and less precisely estimated reduction in support for Lega. At the same time, exposure to immigrant-related crime news increases support for the Center-Right coalition by approximately 2.1 percentage points. These effects indicate a reallocation of votes away from parties with either ambiguous or heterogeneous positions on immigration toward parties that place greater emphasis on law-and-order policies.

4.1. Interpretation and Conceptual Framework

All estimates discussed above reflect within-individual changes in voting behavior across elections, relative to the same individuals past vote choices, rather than cross-sectional differences in baseline preferences.

The patterns documented above indicate that crime-related news does not affect voting behavior mechanically or uniformly. Instead, electoral responses emerge only when crime-related information activates salient political narratives linking crime, immigration, and security.

In this perspective, crime-related news attributed to Italian citizens does not convey new or politically informative content. Such events are likely interpreted as idiosyncratic or non-systemic, and therefore fail to trigger meaningful revisions in vote choice. By contrast, crime-related news attributed to immigrants resonates with a broader and highly politicized narrative that associates immigration with security concerns. This narrative is actively emphasized in political discourse and media framing, particularly by parties on the right.

Our findings align with theories of issue ownership, according to which voters respond to shocks by reallocating support toward parties perceived as more competent or credible on the relevant issue. In the Italian context, law and order and immigration control are core dimensions of political competition on the right. When crime-related news explicitly references immigrants, voters appear to update their beliefs about which parties are better positioned to address these concerns, leading to systematic vote reallocation.

Importantly, the results also highlight asymmetries within the right-wing bloc. Parties with ambiguous or internally heterogeneous positions on immigration—such as the Five Star Movement and, to a lesser extent, Lega—experience electoral losses following exposure to immigrant-related crime news. In contrast, parties with a more traditional and unambiguous law-and-order platform benefit from these shocks. This pattern suggests that ambiguity on highly salient issues may be electorally costly when voters are confronted with information that heightens perceived security risks.

More broadly, the contrast between aggregate crime exposure and nationality-specific crime exposure underscores the importance of framing and attribution. When crime-related news lacks clear attribution, its electoral impact is weak and unstable. Once attribution is made salient, however, voter responses become systematic and directional. This finding contributes to a growing literature emphasizing that political reactions to crime are mediated not only by crime itself, but by how crime is narrated and linked to broader social categories.

Heterogeneity. The conceptual framework developed above implies that electoral responses to crime-related news should not be uniform across voters. If crime exposure operates by activating narratives about immigration and security, responsiveness should vary systematically across groups that differ in political sophistication, economic vulnerability, and baseline exposure to these narratives.

We therefore examine heterogeneity along key demographic and socioeconomic dimensions, focusing in particular on education and skill levels. These dimensions are plausibly related to how individuals interpret and react to crime-related information, as well as to their prior beliefs about immigration and security. By contrast, heterogeneity by parental status and gender does not yield qualitatively different insights and is therefore reported in Appendix C.

This heterogeneity analysis helps clarify which segments of the electorate are most responsive to crime-related news exposure and provides additional evidence in support of the mechanism proposed above.

4.2. Heterogeneity Effects on National Elections

The conceptual framework developed above also yields clear predictions about heterogeneous electoral responses to crime-related news. If crime exposure affects voting behavior primarily by activating narratives linking crime, immigration, and security, responsiveness should vary systematically across voter groups that differ in political sophistication, economic vulnerability, and baseline exposure to these narratives. In this subsection, we investigate heterogeneity along education and occupational skill dimensions (Table 2), which proxy for differences in information processing, political awareness, and sensitivity to law-and-order framing.

Table 2: Differential Effects of Crime Perception on Voting at National Elections

Panel A: By Education Level

Dep.	M5S		Lega		Center Right		Center Left		Abstention	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
Crime by Italians × low education	-0.015	-0.014	0.024	0.021	0.002	0.004	0.008	0.006	-0.005	-0.003
	(0.010)	(0.010)	(0.023)	(0.026)	(0.022)	(0.024)	(0.009)	(0.009)	(0.019)	(0.018)
Crime by Italians × high education	0.002	0.002	-0.013	-0.016	0.002	0.004	0.001	0.000	0.011	0.012
	(0.015)	(0.016)	(0.018)	(0.017)	(0.017)	(0.017)	(0.008)	(0.007)	(0.006)	(0.007)
Crime by immigrants × low education	-0.013	-0.012	-0.039***	-0.034**	0.029*	0.030*	-0.005	-0.003	0.016	0.012
	(0.011)	(0.012)	(0.009)	(0.012)	(0.013)	(0.014)	(0.010)	(0.010)	(0.013)	(0.012)
Crime by immigrants × high education	-0.027**	-0.026**	-0.003	0.002	0.016	0.016	0.013	0.015	-0.001	-0.005
	(0.008)	(0.008)	(0.008)	(0.009)	(0.009)	(0.010)	(0.009)	(0.010)	(0.014)	(0.015)
Individual fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓		✓		✓		✓		✓	
Districts × Year fixed effects		✓		✓		✓		✓		✓
R^2	0.05	0.05	0.16	0.16	0.11	0.11	0.02	0.02	0.01	0.01
Observations	9624	9624	9624	9624	9624	9624	9624	9624	9624	9624

Panel B: By Skills Level

Dep.	M5S		Lega		Center Right		Center Left		Abstention	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
Crime by Italians × low skills	-0.012**	-0.011*	0.009	0.006	-0.011	-0.009	0.017	0.015	0.001	0.003
	(0.004)	(0.004)	(0.016)	(0.018)	(0.016)	(0.017)	(0.011)	(0.011)	(0.008)	(0.008)
Crime by Italians× high skills	0.013	0.014	-0.019	-0.021	0.030	0.032	-0.026	-0.027	0.014*	0.015**
	(0.015)	(0.016)	(0.027)	(0.025)	(0.024)	(0.024)	(0.017)	(0.017)	(0.006)	(0.005)
Crime by immigrants × low skills	-0.007	-0.006	-0.011	-0.006	0.028**	0.028*	0.002	0.003	0.011	0.007
	(0.012)	(0.012)	(0.007)	(0.010)	(0.010)	(0.011)	(0.015)	(0.016)	(0.012)	(0.014)
Crime by immigrants × high skills	-0.048***	-0.046***	-0.022	-0.017	0.008	0.008	0.016	0.019	-0.006	-0.010
	(0.005)	(0.004)	(0.015)	(0.014)	(0.006)	(0.007)	(0.017)	(0.016)	(0.020)	(0.020)
Individual fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓		✓		✓		✓		✓	
Districts × Year fixed effects		✓		✓		✓		✓		✓
R^2	0.05	0.05	0.16	0.16	0.11	0.12	0.02	0.02	0.01	0.01
Observations	9624	9624	9624	9624	9624	9624	9624	9624	9624	9624

Notes: This table reports OLS estimates of the effect of local crime on voting behavior at national elections. The unit of observation is the individual i at time t. The dependent variable is an indicator equal to one if the respondent votes for M5S, Lega, Center Left, Center Right, or abstains from voting. The key explanatory variables are indicators for whether at least one crime is committed by Italians or immigrants within 200 meters of the respondents street in the 30 days preceding the election day. These variables are interacted with individual characteristics capturing heterogeneity by education (Panel A) and by skill level (Panel B). Robust standard errors, clustered at the district level, are reported in parentheses.

Education. We begin by examining heterogeneity by educational attainment. Education plausibly captures voters ability to interpret and contextualize crime-related information within broader political narratives. More educated voters may be better positioned to recognize ideological inconsistencies or ambiguity in parties stances on immigration and security, and therefore more likely to update their vote choice when salient crime-related news conflicts with these positions.

Consistent with this interpretation, the estimates in Tab. 2A indicate that highly educated voters are more likely to abandon the Five Star Movement (M5S) following exposure to immigrant-related crime news. By contrast, lower-educated voters exhibit stronger disengagement from Lega and are more likely to reallocate support toward the Center-Right coalition. This pattern suggests that education shapes not only the intensity of electoral reactions, but also the direction of vote reallocation, with more educated voters penalizing ideological ambiguity and less educated voters responding more directly to law-and-order appeals.

Occupational skill. We next turn to heterogeneity by occupational skill level, which captures differences in labor-market vulnerability and exposure to redistributive and security concerns. Lower-skilled voters may perceive crime-related shocks as more directly threatening to economic and social stability, thereby placing greater weight on parties with clear and traditional law-and-order platforms.

The results in Tab. 2B closely mirror those by education. Highly skilled voters disproportionately exit M5S following immigrant-related crime exposure, while lower-skilled voters exhibit a systematic shift toward the Center-Right coalition. Together, these findings reinforce the interpretation that crime-related news affects voting behavior through the activation of security-related narratives, with heterogeneous responses reflecting differences in political sophistication and perceived exposure to the underlying risks emphasized by these narratives.

4.3. Mechanism: Framing, Issue Ownership, and Selective Responsiveness

The heterogeneity patterns documented above provide important clues about the mechanism through which crime-related news affects voting behavior. Responses to immigrant-related crime are not uniform across voters, but are systematically stronger among groups that differ in political sophistication, occupational position, and baseline alignment with competing narratives on immigration and security.

These structured differences suggest that crime-related news does not operate as a generic information shock. Instead, electoral responses appear to emerge when crime reporting activates salient political frames that link immigration to law and order, thereby interacting with parties perceived issue ownership on these dimensions. In this section, we articulate a conceptual framework that rationalizes the main and heterogeneous effects through a mechanism of framing-driven selective responsiveness.

A pure crime-salience channel would predict broadly similar and monotonic electoral reactions to local crime news regardless of offender characteristics, with effects increasing smoothly in proximity or frequency. Our results do not support this interpretation. Aggregate measures of crime exposure yield weak and unstable effects, and crimes attributed to Italian citizens generate no systematic electoral response, even when they occur in the immediate vicinity of voters.

By contrast, electoral reactions emerge sharply and consistently when crime is explicitly attributed to immigrants. This asymmetry suggests that nationality attribution operates as a framing device that activates politically salient associations between immigration, security, and public order. Importantly, this interpretation does not require voters to update beliefs about actual crime rates, but only to update the perceived relevance of immigration as a political issue in the electoral domain. In this perspective, immigrant-related crime news does not simply convey information about local criminal activity, but resonates with an existing and highly politicized narrative that links immigration to social threat.

The direction of vote reallocation further supports this interpretation. Exposure to immigrant-related crime news induces voters to withdraw support from parties with ambiguous or internally heterogeneous positions on immigration—most notably the Five Star Movement—and to shift toward parties with a clearer law-and-order profile. This pattern is consistent with theories of issue ownership, whereby voters respond to salient shocks by reallocating support toward parties perceived as more competent or credible on the activated issue dimension.

Additional support for a framing-based mechanism comes from the contrast between national and administrative elections. In national contests, immigrant-related crime news leads to systematic vote switching across parties. In local elections, instead, the same type of news primarily increases abstention. This divergence suggests that the mechanism operates through political interpretation rather than through generalized fear or anger: when party switching in response to immigration-related narratives is socially or ideologically costly—as in local elections—voters respond by disengaging rather than by realigning.

Finally, heterogeneity by education and skill levels reinforces this interpretation. More educated and high-skilled voters disproportionately penalize parties perceived as ideologically inconsistent on immigration, while lower-skilled voters shift toward traditional law-and-order parties. These patterns are consistent with selective responsiveness to narrative framing, rather than with uniform reactions to crime risk.

Taken together, the evidence points to a mechanism in which crime-related news affects voting behavior by selectively activating identity-based political narratives, with electoral consequences shaped by party positioning and institutional context. While our data do not allow us to observe belief updating directly, the structure, asymmetry, and contextual dependence of the estimated effects strongly favor a framing-based interpretation over alternative explanations based on generic crime salience or mechanical responses to local crime exposure.

4.4. Administrative Elections

We now turn to administrative elections, which provide a complementary setting to assess how the electoral impact of crime-related news depends on the level of government at which political accountability operates. Compared to national elections, local contests are characterized by stronger incumbency considerations and clearer attribution of responsibility for local conditions, including public safety.

Table 3 reports estimates analogous to those presented for national elections, distinguishing between crime-related news attributed to Italian citizens and to immigrants.

Table 3: Differential Effects of Crime Perception on Voting at Administrative Elections

Dep.	M5S		Center Right & Lega		Center Left		Abstention	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
Crimes by Italians	0.003	0.002	0.005	0.007	-0.015*	-0.015**	0.005	0.003
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Crimes by immigrants	-0.003	-0.002	0.004	-0.000	-0.011	-0.007	0.012**	0.014**
	(0.006)	(0.006)	(0.012)	(0.012)	(0.011)	(0.013)	(0.005)	(0.005)
Individual fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓		✓		✓		✓	
Districts × Year fixed effects		✓		✓		✓		✓
R^2	0.05	0.05	0.27	0.27	0.14	0.14	0.00	0.00
Observations	14436	14436	14436	14436	14436	14436	14436	14436

Notes: This table reports OLS estimates of the effect of local crime on voting behavior at administrative elections. The unit of observation is the individual i at time t. The dependent variable is a dummy that takes the value 1 if the individual votes for parties such as M5S, Lega, Center Left, Center Right, or abstains from voting. The independent variables are dummies that takes value 1 if there is at least one crime committed by Italians and by immigrants, respectively, within 200 meters of the respondents street in the 30 days preceding the election day. Robust standard errors are clustered at district level.

Table 3 reports estimates analogous to those presented for national elections, distinguishing between crime-related news attributed to Italian citizens and to immigrants.

The results in Table 3 reveal a qualitatively different pattern from that observed in national elections. In administrative contests, exposure to crime-related news attributed to Italian citizens leads to a reduction in support for the Center-Left, consistent with electoral punishment of local incumbents. This pattern aligns with standard accountability mechanisms, whereby voters hold local governments responsible for perceived failures in maintaining public order.

By contrast, when crime-related news is attributed to immigrants, we observe an increase in abstention rather than a systematic reallocation of votes across parties. This response suggests that, in local elections, crime narratives linked to immigration do not translate into clear electoral alternatives for voters who are dissatisfied with incumbent performance. Instead of switching parties, some voters disengage from the electoral process altogether.

Taken together, these findings underscore that the political consequences of crime-related news depend not only on attribution, but also on the institutional context in which elections take place. While national elections facilitate vote reallocation across parties competing on immigration and security platforms, local elections emphasize incumbent accountability, leading either to electoral punishment or to abstention when credible alternatives are absent.

5. Identification Concerns

In this section, we provide additional evidence supporting the plausibility of our identifying assumptions and the robustness of our main estimates.

5.1. Crime News Coverage and Political Orientation Around Elections

A key identifying assumption of our empirical strategy is that local crime news coverage is not systematically influenced around elections by the political orientation of neighborhoods. In particular, we assume that neither strategic editorial choices by local media outlets nor differential reporting behavior by residents—such as a higher propensity to report crimes or suspicious activities to authorities in certain politically aligned areas—drive changes in crime news coverage during the electoral period. If this assumption were violated, observed variation in crime-related news could reflect partisan reporting dynamics rather than the arrival of salient local information.

To assess the plausibility of this assumption, we examine whether crime news coverage evolves differently around elections in districts historically aligned with the right and with the left. Rather than estimating a causal effect, this exercise serves as a diagnostic check for potential strategic behavior in news reporting around electoral events.

Specifically, we implement a difference-in-discontinuities style analysis around the 2013 and 2018 national elections. For each week relative to the election date, we compute the difference in crime-related news coverage between right-leaning and left-leaning districts, where political orientation is assigned using vote shares from the 2008 general election, which predates our sample period.

Figure 1 plots the evolution of this right–left difference in local crime news coverage in the 14 weeks before and after the elections. Figure 1a focuses on crime news involving Italian citizens, while Figure 1b considers crime news involving immigrants. In both panels, solid lines represent linear fits estimated separately on each side of the election threshold, with dashed lines indicating 95% confidence intervals.

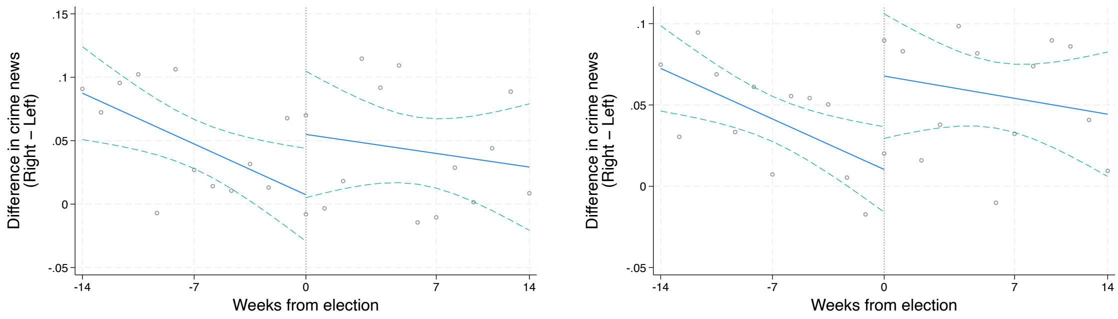

(a) Crime news involving Italian citizens.

(b) Crime news involving immigrants.

Fig. 1: Differences in crime news coverage between right- and left-leaning districts around national elections.

Notes: This figure plots the difference in district-level weekly average shares of crime-related news mentions between right-leaning and left-leaning districts (Right minus Left). The time window spans 14 weeks before and after the 2013 and 2018 national elections. District political orientation is assigned using vote shares from the 2008 general election. Each dot represents the weekly difference in average crime news coverage. Solid lines are linear fits estimated separately on each side of the election threshold, with dashed lines denoting 95% confidence intervals.

Across both types of crime, we find no evidence of a discontinuous change in the right–left difference at the election threshold. Pre- and post-election trends appear smooth and continuous, and the estimated confidence intervals rule out economically meaningful jumps at the cutoff.

This visual evidence is corroborated by formal difference-in-discontinuities tests, reported in Appendix C, which are consistent with the absence of statistically significant jumps at the election threshold.

Taken together, these results provide no indication that local crime news coverage is strategically adjusted around elections in response to district political alignment. This finding supports our interpretation of media-reported crime as reflecting the diffusion of local information rather than short-run political manipulation.

5.2. Testing identifying assumptions: placebo effect of crime news after the elections

In our primary analysis, we assume that crime perception due to news articles influence voting behavior primarily due to their publication during the political campaign period. To validate this assumption, we estimate the effects of crime perception within the 30-day window preceding the elections, while excluding the first two weeks to mitigate potential persistence effects of news coverage both before and after the electoral event. This approach allows us to isolate the immediate impact of crime-related media exposure on voter decisions, ensuring that our results are not driven by longer-term trends or post-election reporting.

In our primary analysis, we assume that crime-related news affects voting behavior primarily when it occurs during the electoral campaign period. To assess this assumption, we conduct placebo tests using crime-related news exposure occurring after elections. If our main results were driven by persistent trends or spurious correlations, we would expect to observe similar effects following election dates.

Accordingly, we estimate specifications analogous to our baseline models using crime-related news exposure within a 30-day window after elections, excluding the first two weeks to mitigate mechanical overlap with pre-election reporting. This approach allows us to isolate post-election exposure and assess whether crime-related news continues to influence voting behavior outside the campaign period.

Table 4: Differential Effects of Crime Perception on Voting after National Elections

Dep.	M5S (1)	M5S (2)	Lega (1)	Lega (2)	Center Right (1)	Center Right (2)	Center Left (1)	Center Left (2)	Abstention (1)	Abstention (2)
Crimes by Italians	-0.010	-0.009	-0.012	-0.011	0.008	0.009	-0.010	-0.008	0.015	0.013
	(0.006)	(0.006)	(0.010)	(0.011)	(0.010)	(0.010)	(0.015)	(0.015)	(0.009)	(0.010)
Crimes by immigrants	-0.001	0.003	0.000	-0.001	-0.004	-0.001	0.005	0.006	-0.001	-0.002
	(0.008)	(0.008)	(0.013)	(0.016)	(0.015)	(0.016)	(0.018)	(0.017)	(0.009)	(0.008)
Individual fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓		✓		✓		✓		✓	
Districts fixed effects × Year fixed effects		✓		✓		✓		✓		✓
R^2	0.05	0.05	0.16	0.16	0.11	0.11	0.02	0.02	0.01	0.02
Observations	9624	9624	9624	9624	9624	9624	9624	9624	9624	9624

Notes: This table reports OLS estimates of the effect of local crime-related news exposure on voting behavior **after** national elections. The unit of observation is the individual i at time t. The dependent variable is a dummy equal to one if the individual votes for parties such as M5S, Lega, Center Left, Center Right, or abstains from voting. The independent variables are indicators equal to one if at least one crime-related news item attributed to Italian citizens or immigrants, respectively, is reported within 200 meters of the respondents street of residence in the 30 days **following** the election day. Robust standard errors are clustered at the district level.

Upon examining the results presented in Table 4, we find no statistically significant effects associated with post-election crime-related news exposure. This suggests that our main estimates are not driven by spurious correlations or persistent trends in crime reporting. The absence of significant post-election effects strengthens the credibility of our identification strategy and supports the interpretation that the observed impact of crime-related news on voting behavior is specific to the electoral campaign period.

6. Conclusion

This paper studies how crime-related news shapes individual voting behavior, focusing on the role of media salience and offender attribution in Italian elections. Using geolocated crime news matched to a panel of non-relocating voters observed across multiple electoral cycles, we examine whether exposure to crime-related information in the immediate pre-election period affects electoral choices.

Our findings reveal a sharp asymmetry. When crime-related news does not specify offenders nationality, electoral responses are weak, unstable, and difficult to interpret. By contrast, when crime is explicitly attributed to immigrants, voters exhibit systematic and directional reactions. In national elections, exposure to immigrant-related crime news leads voters to withdraw support from parties with ambiguous or internally heterogeneous positions on immigrationmost notably the Five Star

Movementand to reallocate votes toward parties with a clear law-and-order profile. Crimes attributed to Italian citizens, instead, do not generate meaningful electoral responses.

These effects are heterogeneous across voter characteristics. Highly educated and high-skilled voters disproportionately abandon the Five Star Movement, while lower-skilled and less-educated voters are more likely to shift away from Lega. This pattern is consistent with the idea that crime-related news affects voting behavior through the activation of salient political narratives linking immigration and security, rather than through generalized concerns about crime.

Electoral responses differ in administrative elections, where accountability and social context are more localized. In this setting, crimes attributed to Italian citizens tend to generate punishment of incumbent administrations, while immigrant-related crime news primarily increases abstention. This suggests that, at the local level, voters may perceive switching parties in response to immigration-related crime as socially or ideologically costly, opting instead for electoral disengagement.

Taken together, these results highlight that crime-related news does not influence voting behavior mechanically. Its electoral impact depends critically on attribution, framing, and the political narratives it activates. By showing that voter responses hinge on how crime is reported rather than on crime per se, this study contributes to the literature on crime perception, media influence, and electoral behavior, and underscores the importance of narrative framing in shaping democratic accountability.

An important consideration for external validity concerns the political context of Bologna, which has historically been characterized by strong center-left dominance and relatively limited electoral support for anti-immigration platforms. In this sense, the estimated effects should be interpreted as a lower bound of the electoral impact of crime-related news in more politically competitive or right-leaning contexts, where immigration and security concerns may be more salient.

Future research could further refine these insights by incorporating richer information on offender–victim pairings and media framing, allowing for a more granular analysis of how different crime narratives shape political behavior.

References

[1] Ajzenman, N., Dominguez, P., and Undurraga, R. (2023). Immigration, crime, and crime (mis)perceptions. *American Economic Journal: Applied Economics*, 15(4):142–176.

[2] Alesina, A., Miano, A., and Stantcheva, S. (2023). Immigration and redistribution. *The Review of Economic Studies*, 90(1):1–39.

[3] Barone, G., D'Acunto, F., and Narciso, G. (2015). Telecracy: Testing for channels of persuasion. *American Economic Journal: Economic Policy*, 7(2):30–60.

[4] Berti Ceroni, C., Piemontese, L., Prarolo, G., and Schiavone, A. (2025). Settlers and seekers: Immigrant proximity and voter polarisation. Working Paper.

[5] Blanes i Vidal, J. and Mastrobuoni, G. (2018). Police patrols and crime. Technical Report 11393, IZA.

[6] Bordalo, P., Coffman, K., Gennaioli, N., and Shleifer, A. (2018). Stereotypes. *The Quarterly Journal of Economics*, 133(4):1753–1794.

[7] Buonanno, P., Montolio, D., and Raya-Vílchez, J. M. (2013). Housing prices and crime perception. *Empirical Economics*, 45(1):305–321.

[8] Couttenier, M., Hatte, S., Thoenig, M., and Vlachos, S. (2024). Anti-muslim voting and media coverage of immigrant crimes. *The Review of Economics and Statistics*, 106(2):576–585.

[9] DellaVigna, S. and Gentzkow, M. (2010). Persuasion: Empirical evidence. *Annual Review of Economics*, 2:643–669.

[10] DellaVigna, S. and Kaplan, E. (2007). The Fox News effect: Media bias and voting. *The Quarterly Journal of Economics*, 122(3):1187–1234.

[11] Drago, F., Galbiati, R., and Sobbrio, F. (2020). The political cost of being soft on crime: Evidence from a natural experiment. *Journal of the European Economic Association*, 18(6):3305–3336.

[12] Durante, R. and Knight, B. (2012). Partisan control, media bias, and viewer responses: Evidence from Berlusconi's Italy. *Journal of the European Economic Association*, 10(3):451–481.

[13] Dustmann, C. and Fasani, F. (2016). The effect of local area crime on mental health. *The Economic Journal*, 126(593):978–1017.

[14] Enikolopov, R., Petrova, M., and Zhuravskaya, E. (2011). Media and political persuasion: Evidence from Russia. *American Economic Review*, 101(7):3253–3285.

[15] Fetzer, T. (2019). Did austerity cause Brexit? *American Economic Review*, 109(11):3849–3886.

[16] Gentzkow, M., Kelly, B., and Taddy, M. (2019). Text as data. *Journal of Economic Literature*, 57(3):535–574.

[17] Gerber, A. S., Karlan, D., and Bergan, D. (2009). Does the media matter? a field experiment measuring the effect of newspapers on voting behavior and political opinions. *American Economic Journal: Applied Economics*, 1(2):35–52.

[18] Grembi, V., Nannicini, T., and Troiano, U. (2016). Do fiscal rules matter? a difference-in-discontinuities design. *American Economic Journal: Applied Economics*, 8(3):1–30.

[19] Guiso, L., Herrera, H., Morelli, M., and Sonno, T. (2024). Economic insecurity and the demand for populism in Europe. *Economica*, 91(362):588–620.

[20] Kahneman, D. and Tversky, A. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2):207–232.

[21] Keita, S., Renault, T., and Valette, J. (2023). The usual suspects: Offender origin, media reporting and natives' attitudes towards immigration. *The Economic Journal*, 134(657):322–362.

[22] Kirchmaier, T., Machin, S. J., and Villa-Llera, C. (2020). Gangs and knife crime in London. Working Paper.

[23] Mastrorocco, N. and Minale, L. (2018). News media and crime perceptions: Evidence from a natural experiment. *Journal of Public Economics*, 165:230–255.

[24] Muço, A. (2025). The politician, the party, and the president: How do political scandals propagate across the party network? *Journal of Economic Behavior & Organization*, 231:106897.

[25] Sacco, V. F. (1982). The effects of mass media on perceptions of crime: A reanalysis of the issues. *The Pacific Sociological Review*, 25(4):475–493.

[26] Vargas, J. F., Purroy, M. E., Coy, F., Perilla, S., and Prem, M. (2025). Fear to vote: Explosions, salience, and elections. Working paper, submitted. Latest version: April 25, 2025.

Appendix

A. Survey Data, Local Context, and Electoral Validation

This Appendix provides background on the local political context, details on the survey design, and validation exercises comparing survey-based voting reports with official electoral outcomes. Together, these elements support the internal consistency of the data and clarify the scope of external validity of our findings.

A.1. Local Political Context: Bologna

Bologna is the seventh most populous Italian city, with nearly 400,000 residents and sustained population growth over the past two decades (+15%). Civic engagement is traditionally high, with voter turnout consistently exceeding the national average. Historically, Bologna has been a stronghold of the center-left, although recent elections have witnessed increasing support for populist parties, particularly the Five Star Movement and, to a lesser extent, the Lega (Fetzer, 2019; Guiso et al., 2024).

Crime and public safety are salient issues in local political debate and receive substantial media attention. Local news outlets frequently report crime-related events at a highly granular geographic level, contributing to sustained public exposure to crime-related information.

This political and informational environment makes Bologna a particularly suitable setting to study the electoral consequences of crime-related news exposure. At the same time, its historically left-leaning electorate implies that estimated effects—especially those related to immigration and law-and-order narratives—should be interpreted as a lower bound relative to contexts where right-wing parties have stronger baseline support. In settings with weaker left-wing dominance or higher political polarization, similar crime-related news shocks may plausibly generate larger electoral responses.

A.2. Survey Design and Sample Construction

The survey was administered in December 2021 to 5,000 Italian citizens who had resided in Bologna continuously since at least 2013. Respondents were stratified across the citys 18 neighborhoods, with 70% interviewed by telephone (C.A.T.I.) and 30% online (C.A.W.I.). The questionnaire collected detailed demographic and socio-economic information, including year and place of birth, education, occupation, and family composition, as well as retrospective voting behavior in national and municipal elections between 2004 and 2021.

A distinctive feature of the survey is the collection of respondents street of residence (excluding house numbers), which allows geo-localization at a fine spatial scale. This feature enables us to match individuals to highly localized measures of exposure to crime-related news and to exploit within-city spatial variation in media coverage.

The empirical analysis focuses on a subsample of 4,812 individuals who did not change residence during the observation period. Restricting attention to non-relocating individuals ensures that changes in crime-related news exposure are not mechanically driven by residential mobility, thereby strengthening the interpretation of within-individual variation in exposure across elections.

To mitigate recall bias in retrospective voting reports, voting questions were asked in chronological order and preceded by neutral reminders of the political context of each election. Importantly, the empirical strategy relies on within-individual changes in reported vote choices across elections. As a result, concerns related to level misreporting are mitigated as long as recall errors are not systematically correlated with localized crime-related news exposure.

Survey mode (C.A.T.I. versus C.A.W.I.) is orthogonal to both election timing and local crime exposure. Moreover, all specifications include individual fixed effects, which absorb any time-invariant differences in reporting behavior across respondents and survey modes.

A.3. Electoral Validation

To assess the external validity of survey-based voting reports, we compare self-reported vote shares with official electoral outcomes. Figure A.1 plots survey-based and official vote shares across the five elections used in the empirical analysis and shows a close correspondence for all major parties. Minor deviations emerge in the 2011 and 2016 municipal elections, when the Center-Right and Lega fielded joint candidates, complicating direct comparisons.

While this validation is conducted at the aggregate level, it provides reassurance that systematic misreporting is limited. This is particularly relevant given that our empirical analysis focuses on within-individual changes in voting behavior rather than on cross-sectional vote levels.

Summary statistics for the final sample are reported in Table A.1. Respondents were on average born in 1970, lived in households of approximately three members, and attained secondary or post-secondary education. Females account for 54% of the sample, and roughly half of respondents were employed at the time of the survey.

Taken together, the descriptive evidence and validation exercises support the suitability of the survey data for analyzing within-city electoral behavior in Bologna and for studying how localized exposure to crime-related news shapes individual voting decisions.

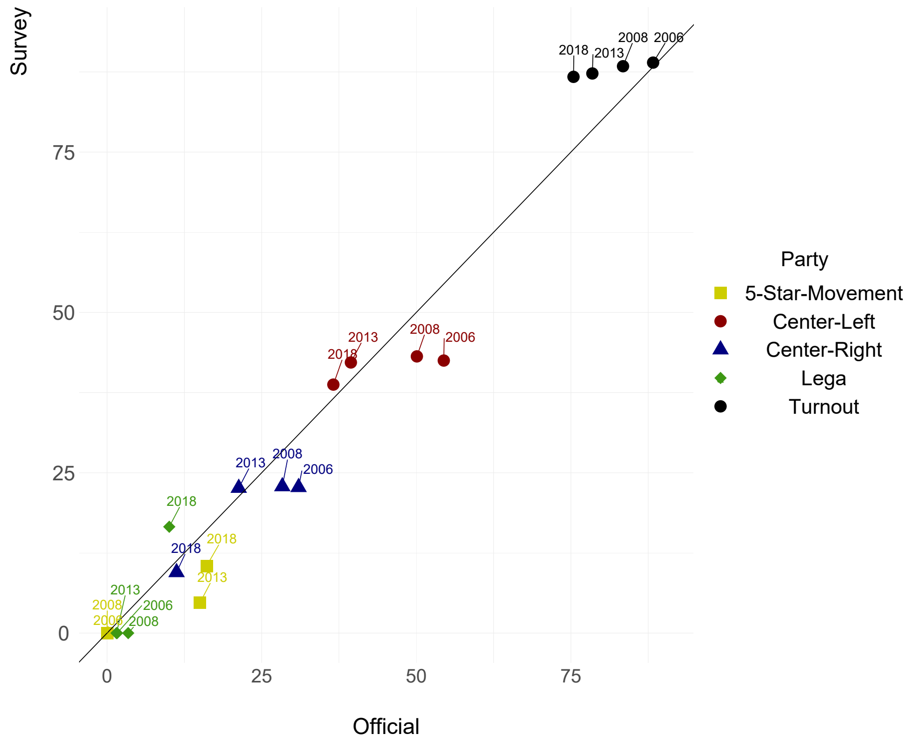

Fig. A.1: Comparison between survey-based and official vote shares

Notes: The figure compares survey-based and official vote shares in the five elections used in the empirical analysis.

Table A.1: Descriptive Statistics of the Survey Sample

Variable	Observations	Mean	Std. Dev.	Min	Max
Individual Characteristics					
Birth Year	4812	1970.145	18.41787	1927	2003
Female	4812	.5367	.4986509	0	1
Household Members	4812	2.954	.922489	1	6
Education (1=PhD,7=Primary School)	4812	5.145	1.109011	1	7
Working	4812	.5054	.4999766	0	1
Student	4812	.1140	.3179048	0	1
Retired	4812	.2640	.4408551	0	1
Homemaker	4812	.0723	.2590022	0	1

Notes: Descriptive statistics refer to the final sample of 4,812 non-relocating individuals observed across nine elections.

B. Local News Corpus on Crime and Geographic Exposure

This Appendix documents the construction of the local crime-related news corpus, the procedures used to classify and geolocate crime-related news events, and the main descriptive patterns of exposure across individuals and neighborhoods.

B.1. Data Source and Collection

Our measures of local crime exposure are based on *BolognaToday*, a leading online local news outlet operating independently of political parties. According to data from *SimilarWeb*, *BolognaToday* attracts approximately 589,000 organic visitors per month. Readers view an average of 7.35 pages per session and spend roughly two minutes per visit. The audience is relatively young and balanced by gender, making the outlet a central source of local information for residents of Bologna.

The archive contains approximately 55,000 articles published between 2010 and 2022. Of these, about 18.5% (roughly 11,000 articles) report on crime-related events. Among crime-related articles, approximately 41% explicitly mention immigrants as perpetrators. Around 75.7% of crime-related articles include sufficiently precise geographic references to allow street-level geolocation, yielding about 7,800 usable observations, of which approximately 46% concern crimes attributed to immigrants.

Importantly, the prominence of crime coverage in *BolognaToday* reflects both the salience of crime in local political debate and editorial choices about which events are deemed newsworthy. As such, our data capture the information environment faced by voters rather than objective crime rates.

Before turning to individual-level exposure measures, we document the main spatial patterns of crime-related news coverage across the city.

B.2. Crime Classification and Textual Processing

Crime-related articles are identified using a dictionary-based classification approach following established methods in computational text analysis (Gentzkow et al., 2019; Muço, 2025). Articles are classified as crime-related when keywords associated with specific offenses (e.g., theft, assault, fraud, rape, murder, kidnapping, and drug-related crimes) co-occur with institutional terms such as "arrest," "reported," or "investigation."

Nationality markers embedded in the text allow us to distinguish between crimes attributed to Italian citizens and those attributed to immigrants. This approach prioritizes transparency and replicability while limiting researcher discretion. Crucially, the classification captures how crime is framed and narrated in the media, rather than the true incidence or severity of criminal activity.

In local crime reporting, references to non-Italian nationality are systematically made explicit, whereas crimes committed by Italian citizens are typically described using place of origin (e.g., city of residence) or initials of the offenders name and surname, without mentioning nationality. As a result, the absence of nationality markers in an article is informative and reflects standard editorial conventions rather than missing information. Accordingly, crime-related articles in which no nationality is specified are classified as involving Italian citizens. This rule reflects a deterministic editorial convention rather than an assumption about offenders characteristics, and therefore avoids differential measurement error across nationality categories. Our analysis therefore captures differences in media salience and framing between crimes attributed to immigrants and those implicitly attributed to Italians.

To assess the reliability of this automated nationality classification, we manually coded a random sample of 100 crime-related articles drawn from the full corpus. Table B.1 reports the resulting confusion matrix. The classifier achieves high accuracy, precision, and recall, indicating that misclassification is limited and unlikely to drive the main results.

Table B.1: Validation of nationality classifier on a manually coded sample

	Predicted Italian	Predicted Immigrant	Total
True Italian	43	5	48
True Immigrant	6	46	52
Total	49	51	100
Accuracy: 0.890; Precision: 0.902; Recall: 0.885; F1: 0.893			

Notes: Confusion matrix comparing automated classification (`immigrants`) to manual coding (`human_imm`) on a random sample of 100 crime-related articles (seed=123456). Rows report manual labels; columns report automated predictions. The composition of the validation sample approximately reflects the distribution of immigrant-related and Italian-related articles in the full corpus.

B.3. Descriptive Spatial Patterns in Crime News Coverage

Crime-related news coverage exhibits pronounced spatial heterogeneity across neighborhoods. Peripheral areas such as *Colli* consistently display no reported crimes, while central districtssuch as *Bolognina*, *Marconi*, and *Irnerio*show persistent concentrations of crime reporting (Figure B.1).

The linguistic content of crime reporting, illustrated in Figure B.2, emphasizes institutional actors (e.g., police, carabinieri), enforcement actions (e.g., arrests), and offense types. Frequent references to nationality and demographic attributes suggest that crime reporting is often framed within identity-based narratives, particularly in articles involving immigrants.

Figure B.3a and Figure B.3b display the spatial distribution of crime-related news reports in the months preceding the 2013 and 2018 elections. Both maps reveal clustering in central districts, with immigrant-related incidents appearing more spatially concentrated. While such clustering may partly reflect underlying patterns of criminal activity, it may also capture editorial selectivity in reporting.

Fig. B.1: Spatial distribution of crime news in the five election years (2011, 2013, 2016, 2018, and 2021) in the month before each election.

Notes: Red dots (●) indicate crimes attributed to immigrants; green dots (●) indicate crimes attributed to Italians.

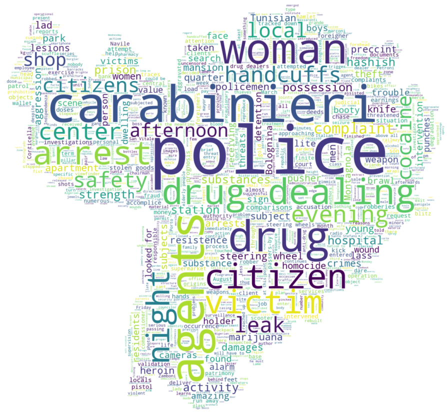

Fig. B.2: Most frequent terms in the crime-related news corpus.

Notes: The figure is a visual representation of text data where words are displayed in varying sizes, with the most frequent words appearing largest and most prominent

(a) 2013.

(b) 2018.

Fig. B.3: Spatial distribution of crime-related news reports by nationality in the month before the given election.

Notes: Red dots (●) indicate crimes attributed to immigrants; green dots (●) indicate crimes attributed to Italians

B.4. Individual Exposure and Measurement Choices

Individual exposure to crime-related news is highly right-skewed. As shown in Figure B.4, most individuals are linked to zero or one crime-related article within the pre-election window, while a small minority experience substantially higher exposure. This skewness is more pronounced for immigrant-related crime news.

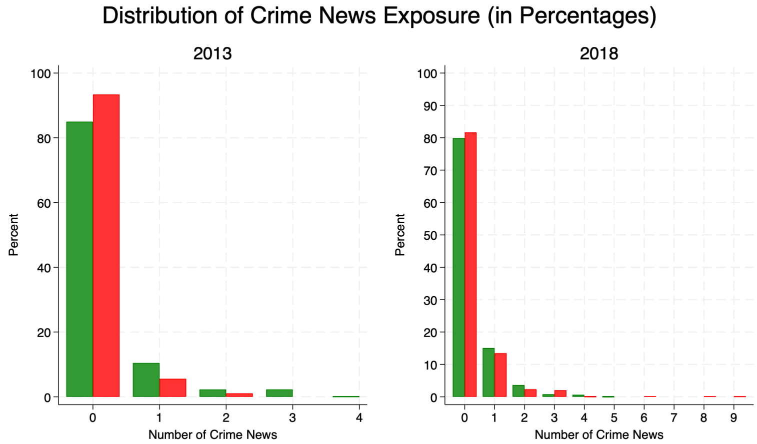

Fig. B.4: Distribution of individual exposure to crime news in 2013 and 2018.

Notes: Exposure is highly right-skewed. Most individuals are linked to zero or one article, while a minority are associated with disproportionately high coverage.

This distribution motivates our use of binary exposure indicators in the main analysis. By focusing on whether an individual is exposed to at least one crime-related news item within a narrow spatial and temporal window, we capture the salience of crime-related information while limiting the influence of extreme observations driven by repeated reporting of the same incident.

Accordingly, our exposure measures should be interpreted as proxies for localized crime salience rather than as measures of crime intensity or frequency. This conservative measurement choice reduces sensitivity to editorial amplification and strengthens the credibility of the estimated electoral effects by ensuring that identification is not driven by repeated reporting of the same incident.

C. Robustness and Additional Results

This Appendix reports robustness checks and additional heterogeneity analyses that complement the baseline results in the main text. Unless otherwise noted, all regressions use the baseline specification and the same sample restrictions as in the main analysis. Across exercises, the qualitative patterns documented in the main text are generally stable: effects attenuate when exposure is defined over broader spatial or temporal windows, consistent with lower salience of more distant or older news.

C.1. Formal test for a discontinuity in crime-news coverage by district ideology

Figure 1 plots the weekly difference in crime-news coverage between right-leaning and left-leaning districts (Right minus Left) around the 2013 and 2018 national elections. To complement the visual evidence, Table C.1 reports a formal discontinuity test.

Specifically, we collapse the data to the election-year × week level and define the outcome as the weekly Right–Left difference in district-level average crime-news shares. We then estimate a piecewise-linear specification that allows slopes to differ on the two sides of the election week and includes election-year fixed effects. The coefficient on *Post-election jump (at week 0)* captures the discontinuity at the threshold.

Consistent with the figure, we do not detect statistically meaningful discontinuities at the election threshold, either for crime news involving Italian citizens or for crime news involving immigrants. Overall, these results do not support the hypothesis that crime-news coverage is strategically adjusted around elections in a way that differentially affects right- versus left-leaning districts.

Table C.1: Difference-in-discontinuities estimates of crime-news coverage around elections

	Italians	Immigrants
Post-election jump (at week 0)	0.396	0.364*
	(0.264)	(0.192)
Weeks from election	-0.061***	-0.034**
	(0.018)	(0.013)
Post × Weeks	0.058*	0.042*
	(0.033)	(0.024)
Election FE	✓	✓
R^2	0.167	0.468
Observations	52	52

Notes: The unit of observation is election-year × week. The dependent variable is the weekly difference in district-level average crime-news shares between right-leaning and left-leaning districts (Right minus Left). The sample includes weeks within ±14 of the 2013 and 2018 national elections. The coefficient on *Post-election jump (at week 0)* captures the discontinuity at the election threshold. Linear trends are allowed to differ before and after the threshold via Post × Weeks. Robust standard errors are reported in parentheses.

C.2. Heterogeneity by districts' historical political orientation

Table C.2 examines whether the electoral response to crime-related news exposure differs by districts historical political orientation. The estimates suggest that reactions to immigrant-related crime news are more pronounced in historically right-leaning districts, where exposure is associated with a lower probability of voting for populist/anti-establishment parties (M5S and, to a lesser extent, Lega) and a corresponding reallocation toward the traditional Center Right. In historically left-leaning districts, estimates are generally smaller and less precisely estimated, although we still observe negative responses for M5S and Lega following immigrant-related crime exposure.

For crime-related news attributed to Italian citizens, effects are comparatively weaker and less systematic across district types, consistent with the main-text evidence that nationality attribution is a key driver of structured electoral responses.

Table C.2: Effects of crime-related news exposure on voting in national elections, by districts' historical political orientation

Dep.	M5S		Lega		Center Left		Center Right		Abstention	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
Crime by Italians in Right districts	0.006**	0.006***	-0.015	-0.018	0.015**	0.012**	0.012	0.015	0.002	0.003
	(0.002)	(0.001)	(0.024)	(0.027)	(0.005)	(0.004)	(0.026)	(0.028)	(0.003)	(0.004)
Crime by Italians in Left districts	-0.010	-0.009	0.008	0.006	-0.002	-0.003	-0.003	-0.002	0.007	0.009
	(0.010)	(0.010)	(0.013)	(0.014)	(0.006)	(0.006)	(0.014)	(0.015)	(0.007)	(0.008)
Crime by immigrants in Right districts	-0.027***	-0.026***	-0.017**	-0.023***	-0.001	-0.001	0.022	0.027*	0.011	0.012
	(0.004)	(0.003)	(0.005)	(0.001)	(0.003)	(0.006)	(0.013)	(0.011)	(0.012)	(0.015)
Crime by immigrants in Left districts	-0.020*	-0.018	-0.015**	-0.002	0.012	0.015	0.020**	0.017	0.001	-0.007
	(0.008)	(0.009)	(0.005)	(0.007)	(0.008)	(0.008)	(0.006)	(0.009)	(0.014)	(0.015)
Individual fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓		✓		✓		✓		✓	
Districts × Year fixed effects		✓		✓		✓		✓		✓
R^2	0.05	0.05	0.16	0.16	0.02	0.02	0.11	0.11	0.01	0.01
Observations	9624	9624	9624	9624	9624	9624	9624	9624	9624	9624

Notes: The unit of observation is individual i in election t. The dependent variable is an indicator equal to one if the respondent votes for the party (or option) reported in the column header, and zero otherwise. Key explanatory variables indicate whether at least one crime-related news item attributed to Italian citizens or to immigrants is reported within 200 meters of the respondents street of residence in the 30 days preceding the election. All specifications include individual fixed effects and election (or district-by-election) fixed effects as in the baseline. Robust standard errors are clustered at the district level.

C.3. Intensity of exposure (article counts)

The baseline analysis uses binary exposure indicators to capture the salience of at least one crime-related news item while limiting sensitivity to repeated reporting of the same incident. As a robustness check, Table C.3 replaces these indicators with the total number of crime-related news items within the same spatial and temporal window. Coefficients can be interpreted as the marginal change in the probability of voting for a given party associated with one additional crime-related news item.

C.4. Alternative distance thresholds

Table C.4 and Table C.5 replicate the baseline analysis using alternative radii. When expanding the radius from 200 meters to 300 meters, estimates are similar in sign and magnitude to the baseline. When using a 500-meter radius, coefficients tend to attenuate and become less precise, consistent with reduced salience of more distant events and with greater measurement noise in broader spatial definitions.

Table C.3: Exposure intensity (article counts)

Dep.	M5S		Lega		Center Right		Center Left		Abstention	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
Crimes by Italians	-0.001	0.000	0.002	-0.001	-0.003	-0.001	-0.002	0.006^{**}	0.005^{*}	-0.000
	(0.004)	(0.004)	(0.008)	(0.009)	(0.008)	(0.008)	(0.002)	(0.002)	(0.004)	(0.004)
Crimes by immigrants	-0.012^{**}	-0.011^{**}	-0.013^{**}	-0.011	0.013	0.014	0.002	0.003	0.006^{**}	-0.001
	(0.004)	(0.003)	(0.005)	(0.007)	(0.006)	(0.008)	(0.002)	(0.002)	(0.005)	(0.004)
Individual fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓		✓		✓		✓		✓	
Districts × Year fixed effects		✓		✓		✓		✓		✓
R^2	0.05	0.05	0.16	0.16	0.11	0.11	0.02	0.02	0.01	0.01
Observations	9624	9624	9624	9624	9624	9624	9624	9624	9624	9624

Notes: Estimates analogous to the baseline specification but replacing the exposure dummies with the number of crime-related news items in the relevant window. Robust standard errors are clustered at the district level.

Table C.4: Alternative distance threshold: 300m radius

Dep.	M5S		Lega		Center Right		Center Left		Abstention	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
Crimes by Italians	-0.013	-0.012	-0.007	-0.010	0.011	0.014	0.010	0.008	0.001	0.003
	(0.011)	(0.012)	(0.011)	(0.010)	(0.009)	(0.008)	(0.008)	(0.008)	(0.005)	(0.004)
Crimes by immigrants	-0.023^{***}	-0.022^{***}	-0.005	0.000	0.012	0.013	0.004	0.005	0.005	0.002
	(0.005)	(0.005)	(0.007)	(0.007)	(0.007)	(0.008)	(0.007)	(0.008)	(0.006)	(0.005)
Individual fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓		✓		✓		✓		✓	
Districts fixed effects × Year fixed effects		✓		✓		✓		✓		✓
R^2	0.05	0.05	0.16	0.16	0.11	0.11	0.02	0.02	0.01	0.01
Observations	9624	9624	9624	9624	9624	9624	9624	9624	9624	9624

Notes: Baseline specification with exposure defined within a 300m radius. Robust standard errors are clustered at the district level.

Table C.5: Alternative distance threshold: 500m radius

Dep.	M5S		Lega		Center Right		Center Left		Abstention	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
Crimes by Italians	-0.010	-0.009	-0.005	-0.008	0.007	0.010	0.002	0.000	0.004	0.007
	(0.008)	(0.008)	(0.007)	(0.006)	(0.011)	(0.011)	(0.006)	(0.007)	(0.008)	(0.008)
Crimes by immigrants	-0.016^{*}	-0.014	-0.009	-0.003	0.004	0.007	-0.000	0.001	-0.004	-0.011^{***}
	(0.006)	(0.008)	(0.010)	(0.010)	(0.008)	(0.011)	(0.004)	(0.005)	(0.005)	(0.003)
Individual fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓		✓		✓		✓		✓	
Districts fixed effects × Year fixed effects		✓		✓		✓		✓		✓
R^2	0.05	0.05	0.16	0.16	0.11	0.11	0.02	0.02	0.01	0.02
Observations	9624	9624	9624	9624	9624	9624	9624	9624	9624	9624

Notes: Baseline specification with exposure defined within a 500m radius. Robust standard errors are clustered at the district level.

C.5. Alternative time windows

Table C.6 extends the exposure window from 30 to 90 days. As expected, point estimates tend to be smaller when older news is included, consistent with decaying salience over time. The qualitative pattern, however, remains aligned with the baseline.

Table C.6: Alternative time window: 200m radius, 90-day exposure

Dep.	M5S		Lega		Center Right		Center Left		Abstention	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
Crimes by Italians	-0.006	-0.003	-0.009	-0.006	0.012	0.013	0.002	0.001	0.007	0.005
	(0.005)	(0.006)	(0.010)	(0.010)	(0.012)	(0.012)	(0.007)	(0.007)	(0.006)	(0.006)
Crimes by immigrants	0.000	-0.002	-0.013**	-0.013*	0.020	0.019	-0.001	-0.004	-0.002	-0.002
	(0.007)	(0.007)	(0.004)	(0.006)	(0.012)	(0.013)	(0.009)	(0.009)	(0.008)	(0.008)
Individual fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓		✓		✓		✓		✓	
Districts fixed effects × Year fixed effects		✓		✓		✓		✓		✓
R^2	0.05	0.05	0.16	0.16	0.11	0.12	0.02	0.02	0.01	0.01
Observations	9624	9624	9624	9624	9624	9624	9624	9624	9624	9624

Notes: Baseline specification with exposure measured over the 90 days preceding the election. Robust standard errors are clustered at the district level.

C.6. Alternative exposure measures: minimum distance

Table C.7 measures exposure using the minimum distance from the respondents street to the closest geolocated crime-related news item in the pre-election window. Results are broadly consistent with the baseline in terms of signs and relative patterns, supporting robustness to alternative exposure metrics.

Table C.7: Alternative exposure metric: minimum distance (30-day window)

Dep.	M5S		Lega		Center Right		Center Left		Abstention	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
Closest crime by Italians	-0.000	0.001	0.012	0.015	-0.015	-0.016	-0.007	-0.004	0.005	0.002
	(0.005)	(0.005)	(0.009)	(0.009)	(0.010)	(0.011)	(0.007)	(0.007)	(0.004)	(0.004)
Closest crime by immigrants	0.007**	0.004	0.011**	0.016	-0.000	-0.005	0.009*	0.025**	-0.003	-0.001
	(0.002)	(0.005)	(0.004)	(0.010)	(0.002)	(0.007)	(0.004)	(0.008)	(0.003)	(0.003)
Individual fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓		✓		✓		✓		✓	
Districts × Year fixed effects		✓		✓		✓		✓		✓
R^2	0.05	0.05	0.16	0.16	0.11	0.11	0.02	0.02	0.01	0.01
Observations	9620	9620	9620	9620	9620	9620	9620	9620	9620	9620

Notes: Baseline specification with exposure measured as minimum distance to the closest crime-related news item in the 30 days preceding the election. Robust standard errors are clustered at the district level.

C.7. Additional heterogeneity analyses

We further explore heterogeneity by parental status and gender. Table C.8 reports estimates for the 30-day window; Table C.9 replicates the analysis using the 90-day window. Overall, these splits do not overturn the main-text interpretation; if anything, they suggest that responsiveness is somewhat stronger among parents, while gender differences are comparatively modest.

Table C.8: Additional heterogeneity: parental status and gender (30-day window)

Panel A: By Having Children

Dep.	M5S (1)	M5S (2)	Lega (1)	Lega (2)	Center Right (1)	Center Right (2)	Center Left (1)	Center Left (2)	Abstention (1)	Abstention (2)
Crime by Italians × w/ children	-0.009*	-0.008	0.005	0.002	0.002	0.004	-0.001	-0.003	0.005	0.007
	(0.005)	(0.005)	(0.007)	(0.008)	(0.013)	(0.014)	(0.008)	(0.009)	(0.008)	(0.008)
Crime by Italians × w/o children	0.010	0.010	-0.010	-0.012	0.002	0.003	0.018	0.017	0.005	0.006
	(0.020)	(0.020)	(0.039)	(0.039)	(0.035)	(0.035)	(0.026)	(0.025)	(0.023)	(0.024)
Crime by immigrants × w/ children	-0.022**	-0.020*	-0.031**	-0.025	0.019	0.019	-0.007	-0.005	0.000	-0.004
	(0.008)	(0.008)	(0.011)	(0.013)	(0.010)	(0.013)	(0.007)	(0.007)	(0.013)	(0.014)
Crime by immigrants × w/o children	-0.024	-0.023	0.031	0.033	0.024	0.025	0.049*	0.048*	0.017	0.014
	(0.015)	(0.014)	(0.030)	(0.030)	(0.026)	(0.024)	(0.019)	(0.020)	(0.012)	(0.010)
Individual fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓		✓		✓		✓		✓	
Districts × Year fixed effects		✓		✓		✓		✓		✓
R^2	0.05	0.05	0.16	0.16	0.11	0.11	0.02	0.02	0.01	0.01
Observations	9624	9624	9624	9624	9624	9624	9624	9624	9624	9624

Panel B: By Gender

Dep.	M5S (1)	M5S (2)	Lega (1)	Lega (2)	Center Right (1)	Center Right (2)	Center Left (1)	Center Left (2)	Abstention (1)	Abstention (2)
Crime by Italians × female	-0.002	-0.001	-0.008	-0.010	0.006	0.008	0.017	0.015	0.001	0.003
	(0.007)	(0.007)	(0.022)	(0.023)	(0.015)	(0.016)	(0.009)	(0.009)	(0.013)	(0.013)
Crime by Italians × male	-0.006	-0.006	0.008	0.006	-0.002	-0.000	-0.009	-0.011	0.009	0.011*
	(0.009)	(0.009)	(0.007)	(0.008)	(0.012)	(0.013)	(0.009)	(0.009)	(0.004)	(0.005)
Crime by immigrants × female	-0.028*	-0.026	-0.015	-0.010	0.020**	0.020*	-0.002	-0.000	0.012	0.008
	(0.013)	(0.013)	(0.011)	(0.011)	(0.007)	(0.008)	(0.016)	(0.016)	(0.011)	(0.011)
Crime by immigrants × male	-0.016*	-0.015*	-0.016	-0.011	0.021	0.021	0.017	0.019	-0.004	-0.008
	(0.007)	(0.007)	(0.016)	(0.018)	(0.012)	(0.013)	(0.009)	(0.010)	(0.012)	(0.012)
Individual fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓		✓		✓		✓		✓	
Districts × Year fixed effects		✓		✓		✓		✓		✓
R^2	0.05	0.05	0.16	0.16	0.11	0.11	0.02	0.02	0.01	0.01
Observations	9624	9624	9624	9624	9624	9624	9624	9624	9624	9624

Notes: Baseline specification estimated separately by subgroup. Exposure is defined within 200m in the 30 days preceding the election. Robust standard errors are clustered at the district level.

Table C.9: Additional heterogeneity: parental status and gender (90-day window)

Panel A: By Having Children

Dep.	M5S		Lega		Center Right		Center Left		Abstention	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
Crime by Italians × w/ children	-0.000	0.002	-0.011	-0.009	0.011	0.012	-0.001	-0.002	0.001	0.000
	(0.006)	(0.007)	(0.013)	(0.013)	(0.017)	(0.016)	(0.009)	(0.008)	(0.007)	(0.007)
Crime by Italians × w/o children	-0.021**	-0.019**	-0.001	0.002	0.013	0.013	0.023	0.012	0.011	0.022
	(0.005)	(0.006)	(0.009)	(0.008)	(0.017)	(0.018)	(0.009)	(0.010)	(0.016)	(0.015)
Crime by immigrants × w/ children	-0.000	-0.002	-0.024**	-0.023**	0.019	0.018	-0.008	-0.011**	-0.006	-0.006
	(0.009)	(0.009)	(0.008)	(0.008)	(0.016)	(0.017)	(0.005)	(0.004)	(0.011)	(0.012)
Crime by immigrants × w/o children	0.001	-0.003	0.021	0.017	0.024	0.023	0.021	0.016	0.008	0.010
	(0.015)	(0.014)	(0.030)	(0.031)	(0.022)	(0.022)	(0.027)	(0.028)	(0.011)	(0.011)
Individual fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓		✓		✓		✓		✓	
Districts × Year fixed effects		✓		✓		✓		✓		✓
R^2	0.05	0.05	0.16	0.16	0.11	0.12	0.02	0.02	0.01	0.02
Observations	9624	9624	9624	9624	9624	9624	9624	9624	9624	9624

Panel B: By Gender

Dep.	M5S		Lega		Center Right		Center Left		Abstention	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
Crime by Italians × female	-0.006	-0.004	-0.006	-0.003	0.004	0.004	0.013	0.012	0.002	0.001
	(0.006)	(0.006)	(0.014)	(0.015)	(0.019)	(0.019)	(0.011)	(0.011)	(0.009)	(0.009)
Crime by Italians × female	-0.005	-0.003	-0.013	-0.010	0.020**	0.021**	-0.010	-0.011	0.011	0.010
	(0.005)	(0.006)	(0.008)	(0.008)	(0.007)	(0.007)	(0.012)	(0.012)	(0.008)	(0.008)
Crime by immigrants × female	0.006	0.004	-0.016	-0.016	0.028*	0.027*	0.002	-0.002	-0.006	-0.005
	(0.012)	(0.012)	(0.010)	(0.010)	(0.012)	(0.013)	(0.005)	(0.005)	(0.009)	(0.010)
Crime by immigrants × female	-0.007	-0.009	-0.009	-0.010	0.011	0.010	-0.004	-0.008	0.001	0.002
	(0.006)	(0.006)	(0.017)	(0.018)	(0.016)	(0.017)	(0.016)	(0.016)	(0.007)	(0.008)
Individual fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓		✓		✓		✓		✓	
Districts × Year fixed effects		✓		✓		✓		✓		✓
R^2	0.05	0.05	0.16	0.16	0.11	0.12	0.02	0.02	0.01	0.01
Observations	9624	9624	9624	9624	9624	9624	9624	9624	9624	9624

Notes: Baseline specification estimated separately by subgroup. Exposure is defined within 200m in the 90 days preceding the election. Robust standard errors are clustered at the district level.

Overall, the robustness exercises in this Appendix indicate that the main qualitative patterns are not an artifact of a specific exposure definition or window, and they behave in the expected direction when salience is mechanically reduced (broader radii or longer windows).

Parenthood, Age, and the Opportunity Cost of Voting: Evidence from Administrative Voter Records*

Marco Rosso†

Abstract

This paper studies how parenthood and parental age are associated with voter turnout using a comprehensive administrative panel covering the universe of registered voters in Bologna across four municipal and national elections between 2004 and 2013. By linking individual turnout records to demographic, fiscal, and residential information, we identify parents, track the age of their children, and follow the same individuals over time. We estimate linear probability models with individual and election-year fixed effects, exploiting within-individual variation to account for permanent differences in civic engagement. On average, parenthood is not associated with lower turnout once individual fixed effects are included. However, substantial heterogeneity emerges over the parental life cycle. Parents of young children vote significantly less than comparable non-parents at younger ages: those with children aged 0–2 and 3–5 exhibit turnout penalties of approximately three to five percentage points. These gaps decline steadily—by about 0.2 percentage points per additional year of parental age—and disappear by around age forty. Parents of older children display no turnout deficit. The participation gap is driven almost entirely by mothers, while fathers turnout remains unaffected. The results are robust to alternative specifications and to controls for residential mobility, neighborhood characteristics, and distance to polling stations. Taken together, the findings highlight the importance of life-cycle factors in shaping political participation and suggest that periods of intensive childcare are associated with temporarily lower electoral engagement. More broadly, the analysis points to a channel through which demographic trends, such as delayed fertility, may have implications for democratic representation.

Keywords: parenthood; age; voter turnout; opportunity cost; administrative data.

JEL: D72; J13; J22.

---

*I greatly benefited from feedback from Tommy E. Murphy, Juan F. Vargas, and Paolo Vanin. I am grateful to Giorgio Bellettini and Carlotta Berti Ceroni for granting access to the administrative turnout data. I gratefully acknowledge financial support from the Italian Ministry of University (PRIN project number 2022SHHM2A). All errors are my own. The pronoun "we" is used throughout the paper for convenience.

†Department of Economics, University of Bologna (email: marco.rosso4@unibo.it).

1. Introduction

Voter turnout is a canonical indicator of democratic health, yet participation has declined across many advanced democracies despite universal suffrage and rising educational attainment. A prominent explanation attributes this trend to generational replacement: cohorts socialized in periods of economic security and individualism exhibit weaker civic norms than those formed under conditions of hardship, leading aggregate turnout to fall as older voters are replaced by younger ones (Smets and van Ham, 2013). At the same time, industrialized countries are undergoing profound demographic change. Fertility rates have dropped sharply, the age at first birth has increased, and populations are aging, raising dependency ratios and concerns about inter-generational equity (Bailey, 2009; Sommer, 2018).

These demographic trends matter politically because family status is closely linked to preferences over public goods such as education, childcare, housing, and welfare provision. If individuals are systematically underrepresented during peak family-formation years, political outcomes may reflect the preferences of demographic groups facing fewer caregiving constraints. In low-fertility contexts such as Italy, where childbearing increasingly occurs after age 30, these dynamics raise questions not only about who votes, but also about whose interests are represented in democratic processes.

Classic theories of political participation emphasize life-cycle dynamics: engagement rises with age, peaks in midlife, and later declines (Wolfinger and Rosenstone, 1980). Marriage and family formation are central transitions within this cycle. Early work showed that marriage stabilizes civic engagement and reduces turnout gaps within households (Stoker and Jennings, 1995), while subsequent research highlighted strong gender asymmetries. A large literature documents that womens political participation is particularly sensitive to resource constraints, social norms, and caregiving responsibilities (Schlozman et al., 1994; Verba et al., 1997; Iversen and Rosenbluth, 2006; Dassonneville and Kostelka, 2021). Recent register-based and survey evidence consistently shows that childbirth reduces womens turnout but has little or no effect on men (Bhatti et al., 2019; Quaranta and Dotti Sani, 2018; Grechyna, 2023). Using Italian administrative data similar in spirit to ours, Bellettini et al. (2023) document that childbirth sharply reduces womens turnout while leaving mens participation largely unaffected.

Despite these advances, two limitations remain in the existing literature. First, most studies treat parenthood as a binary condition, comparing parents to non-parents without distinguishing between parents of infants, preschoolers, school-age children, or adolescents. Yet childcare demands—and thus the opportunity cost of voting—vary sharply over the childs life cycle. Second, parental age and childrens age are tightly correlated: parents of infants tend to be younger, while parents of teenagers are typically in midlife. Disentangling these dimensions is essential to understand whether observed participation gaps reflect the timing of family formation, the stage of child-rearing, or both.

Empirically, addressing these issues is challenging. Much of the existing evidence relies on cross-sectional data or self-reported turnout, which is known to substantially overstate actual participation and cannot track within-person changes over time (Ansolabehere and Hersh, 2012). As a result, it is difficult to separate life-cycle dynamics from persistent differences in civic engagement across individuals.

This paper addresses these gaps using a novel administrative panel from the city of Bologna that links the universe of registered voters across four elections—the 2004 and 2009 municipal elections and the 2008 and 2013 national parliamentary elections—to demographic, fiscal, and residential information. Because voter registration is automatic and polling stations are assigned mechanically based on residence, the data cover all eligible citizens rather than a self-selected sample. We link turnout records to household rosters to identify cohabiting children and their ages, constructing indicators for having any minor child and for five mutually exclusive stages based on the age of the youngest child: nursery (0–2), preschool (3–5), primary school (6–10), middle school (11–13), and high school (14–17). The resulting panel comprises over 500,000 person–year observations for adults aged 18–65.

Our empirical strategy exploits the longitudinal structure of the data. We estimate linear probability models with individual and election-year fixed effects, absorbing time-invariant differences in civic engagement and common electoral shocks. By interacting parental status and childrens life stages with parental age, we explicitly model how the turnout gap evolves over the life cycle. This approach allows us to isolate within-individual changes in participation associated with the timing of parenthood and the progression of children through different stages of dependency. This design follows a growing literature in political science that exploits within-individual variation to study life-course transitions and household dynamics in political participation (Stoker and Jennings, 1995; Dahlgaard, 2018; Dehdari et al., 2022).

The analysis yields three main findings. First, once individual fixed effects are included, there is no average turnout difference between parents and non-parents, indicating that simple comparisons mask substantial heterogeneity. Second, the turnout gap is strongly age-dependent and concentrated among parents of very young children. Parents of infants and preschoolers vote three to five percentage points less than comparable non-parents at younger ages, but this penalty declines steadily—by roughly 0.2 percentage points per additional year of parental age—and vanishes by around age 40. Third, this participation gap is driven almost entirely by mothers; fathers turnout is unaffected.

These findings contribute to the literature in three ways. First, they refine life-cycle models of political participation by showing that the timing and intensity of childcare responsibilities—rather than parenthood per se—shape electoral engagement. Second, they provide new evidence on gendered participation gaps by quantifying how motherhood depresses turnout at specific life stages while fatherhood does not. Third, they link micro-level family dynamics to broader demographic trends, highlighting a channel through which delayed fertility and population aging may affect patterns of democratic participation.

The remainder of the paper proceeds as follows. Section 2 introduces the conceptual framework and empirical predictions. Section 3 describes the institutional background and data. Section 4 outlines the empirical strategy and identification. Section 5 presents the main results, and Section 6 discusses their implications. Additional background material, descriptive statistics, and robustness checks are reported in the Appendix.

2. Conceptual framework and empirical predictions

Here, we outline the mechanisms through which parenthood, parental age, and the life stage of children shape electoral participation. Building on standard rational-choice models of voting, individuals compare the expected benefits of participation with its costs, including the opportunity cost of time.[1] In this framework, parenthood affects turnout primarily because it changes how time is allocated and how binding time constraints are across the life cycle, rather than because it fundamentally alters intrinsic civic norms. Three channels are particularly relevant in the context we study.

Time constraints and the opportunity cost of voting. The arrival of children, and especially very young children, increases the demand for non-delegable caregiving time. Parents must coordinate childcare, work, and household responsibilities, leaving less slack time around election day. Even in a setting like Bologna, where formal voting costs are low thanks to automatic registration and nearby polling stations, the opportunity cost of leaving home, queuing, and casting a ballot can be substantial when childcare demands are intensive. These constraints are likely to be most severe when children are infants or preschoolers and to decline as children become more autonomous. Unequal intra-household specialization implies that these time costs fall disproportionately on mothers, who often bear a larger share of caregiving and domestic work. This channel implies that (i) parents

[1] See, for example, Downs (1957) and the subsequent rational-choice literature on turnout.

of infants and preschoolers should vote less than comparable non-parents, (ii) the turnout penalty should be larger for mothers than for fathers, and (iii) the penalty should attenuate as parents age and as children enter school.

Shifting policy priorities and perceived stakes. Parenthood also changes preferences over public policies, particularly for goods such as childcare, early education, housing, and family-friendly services. When children are very young, these issues are salient but the perceived responsiveness of the political system may be limited if parents do not expect short-run policy changes in opening hours, availability of slots, or flexibility of services. In such circumstances, the expected benefits of voting may not rise enough to offset higher time costs. As children age, parents' policy priorities shift towards schooling and long-run opportunities, which are more visibly linked to government programs and may be perceived as more responsive to electoral pressure. This channel is consistent with a temporary participation gap during early childhood that gradually disappears as children progress through the education system.

Social networks, civic routines, and life-course dynamics. Parenthood reshapes social networks and daily routines. New parents may spend less time in workplaces and civic associations that facilitate political discussion and mobilization, and more time in settings organized around childcare. At the same time, parental age correlates with life-cycle factors that typically raise participation, such as income stability, home-ownership, and stronger local ties. The interaction of these forces implies that the negative effect of intensive caregiving on turnout is superimposed on a positive age gradient in political engagement. As a result, one should not expect a permanent turnout gap between parents and non-parents, but rather a temporary dip in participation during the most demanding childcare stages, followed by convergence—or even catch-up—at later ages.

Taken together, these mechanisms generate the empirical predictions that guide our analysis. First, once permanent individual heterogeneity is absorbed, average differences between parents and non-parents may be small, because the turnout penalty is concentrated in a relatively narrow window of the life cycle. Second, within individuals, the transition into parenthood when children are very young should be associated with a temporary reduction in turnout that diminishes as parents age and as children move into school-age stages. Third, these patterns should be strongly gendered, with larger and more persistent penalties among mothers than among fathers, reflecting unequal childcare responsibilities within households. In the rest of the paper, we test these predictions by exploiting administrative panel data that allow us to track the same individuals across elections as their parental status and children's ages evolve.

These predictions map directly into the empirical specifications estimated below, which exploit within-individual variation to isolate temporary life-cycle deviations from long-run participation patterns.

3. Institutional Background and Data

Our analysis combines administrative voter–turnout records with population registry information for the municipality of Bologna. Italy features automatic voter registration: eligible citizens are enrolled by default and assigned to a polling station based on their registered residential address. As a result, the data do not suffer from self-selection into voter rolls. The core outcome is individual turnout, obtained from the municipal electoral office and digitized from paper attendance registers. We observe whether each eligible voter cast a ballot in four elections held during the study period: the municipal elections of 2004 and 2009 and the national parliamentary elections of 2008 and 2013.

Using anonymized but stable individual identifiers, electoral records are merged to socio-demographic and fiscal archives. The resulting dataset is an unbalanced individual panel with up to four observations per person. Attrition is limited and primarily reflects migration out of the municipality and mortality rather than selective exit from the voter rolls.

We enrich the turnout panel with information from the civil registry and household rosters. The registry provides date of birth, gender, marital status, and individual taxable income, as well as household identifiers and information on individuals position within the household. These data allow us to reconstruct household composition at each election year and to identify cohabiting children.

Because household rosters do not explicitly record parent–child links for all members, we implement a transparent and conservative matching procedure within each household-by-year to identify plausible parent–child relationships. We first classify adults aged 18–65 as *parent-like* if they are household heads or spouses/partners. In households where the head is elderly (aged 65 or above), we additionally classify the middle generation flagged as "child" in the household records as parent-like. Candidate children are defined as household members flagged as "child" (or "grandchild" in elderly-head households) and aged 25 or below.

Within each household-year, we match parent-like adults to candidate children and retain only plausible pairs based on an age-gap restriction: the adult must be older than the child by 15 to 50 years. For each child, we retain up to two candidate parents with the smallest absolute age difference. We then aggregate matches at the adult level, computing the number of matched children and the age of the youngest matched child in the household-year.

While this procedure cannot perfectly reconstruct biological parenthood in all household configurations, any residual misclassification is likely to be classical and to attenuate estimated differences between parents and non-parents rather than generate spurious gradients. This feature is particularly relevant for the interaction specifications used in the empirical analysis, which rely on within-individual changes over time rather than cross-sectional comparisons.

Using the resulting aggregates, we define parental status and construct mutually exclusive indicators based on the age of the youngest matched child: nursery (0–2), preschool (3–5), primary school (6–10), middle school (11–13), and high school (14–17). Adults without matched children are coded as non-parents.

We complement the administrative panel with geocoded information on polling stations. Because polling locations are assigned mechanically based on registered residence, we compute the straight-line distance (in kilometers) between each individuals residence and their assigned polling station. This variable captures a salient logistical component of voting costs. We also construct an indicator for residential mobility between elections. Additional covariates used in the empirical analysis include real (deflated) income and marital status. All specifications include individual fixed effects and election-year fixed effects.

Using administrative turnout records avoids survey misreporting, which is known to substantially overstate participation (Ansolabehere and Hersh, 2012), and aligns with a growing literature exploiting universe-level Italian electoral data to study voting behavior (Cantoni, 2020; Cantoni et al., 2021; Schafer et al., 2022).

3.1. Descriptive statistics and patterns

Table 1 reports descriptive statistics for the restricted analysis sample, comparing adults with and without cohabiting minor children. Parents are substantially more likely to be married, while income distributions are broadly comparable across groups. Average distance to the assigned polling station and residential mobility rates differ only modestly between parents and non-parents.

Figure 1 plots average turnout by age and parental status in the restricted sample. Turnout increases with age for all individuals, while parents exhibit lower participation at younger ages. The gap narrows progressively and becomes negligible at older ages.

These patterns are purely descriptive and do not exploit the longitudinal structure of the data. They are intended to motivate the within-individual analysis that follows rather than to identify causal effects. In particular, the age profiles shown

Table 1: Descriptive statistics by parental status (restricted sample)

	w/o children					w/ children				
	mean	sd	min	max	count	mean	sd	min	max	count
Female	0.492	0.500	0.000	1.000	527876	0.555	0.497	0.000	1.000	194917
Age	42.414	10.108	26.000	59.000	527876	42.324	6.790	26.000	59.000	194917
Real income ()	24557.785	38963.655	0.000	8616915.000	436469	29256.635	43857.445	0.000	6039291.000	171262
Distance from the polling station	4.826	4.183	0.002	77.727	527174	5.053	4.548	0.002	75.330	194686
Moved between elections	0.099	0.298	0.000	1.000	527876	0.110	0.313	0.000	1.000	194917
Married	0.379	0.485	0.000	1.000	527876	0.868	0.338	0.000	1.000	194917
Municipal election	0.501	0.500	0.000	1.000	527876	0.494	0.500	0.000	1.000	194917

Notes: This table reports descriptive statistics for the restricted sample used in the main analysis. "w/children" is defined by the presence of at least one matched cohabiting minor (aged 0–17) in the household-year. Income is deflated to constant euros. Distance is measured in kilometers from the individual's residence to the assigned polling station.

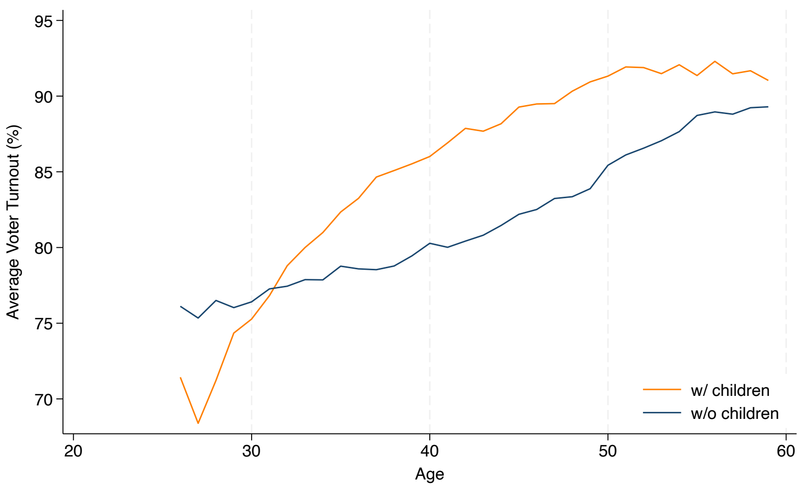

Fig. 1: Average voter turnout by age and parental status (restricted sample)

Notes: The figure plots raw average turnout by age separately for adults with and without matched cohabiting minor children. Turnout is measured as the share of individuals voting in a given election-year. The figure is based on unadjusted means in the restricted sample.

in the figure conflate differences in parental status, childrens age, and parental age, which the empirical strategy explicitly disentangles.

Importantly, the analysis focuses on participation outcomes rather than policy choices. We do not estimate political responsiveness or policy effects, and any implications for democratic representation and accountability should therefore be interpreted as suggestive rather than as estimated causal effects.

For brevity, additional descriptive statistics and figures for alternative subsamples (full sample, women only, and mothers of nursery or preschool-age children) are reported in Appendix B.

4. Empirical Strategy and Identification

This paper exploits a longitudinal administrative panel of voters in Bologna to study how parental status and parental age shape electoral participation relative to non-parents. The dependent variable, $voted_{it}$, is an indicator equal to one if individual i casts a ballot in election year t. The key explanatory variables capture parental status and childrens life stage. In the baseline specification, has_minor_{it} equals one if individual i has at least one cohabiting child under the age of 18 in year t. In more flexible specifications, we distinguish parents by the age of the youngest child, allowing effects to vary across nursery (0–2), preschool (3–5), primary school (6–10), lower secondary (11–13), and upper secondary (14–17) stages.

The panel structure of the data allows us to compare the same individuals across elections and to control for time-invariant unobserved heterogeneity. All specifications include individual fixed effects, α_i, and election-year fixed effects, λ_t, absorbing persistent differences in civic engagement across individuals and common shocks across elections. This approach is consistent with recent studies exploiting within-individual variation to study life-course transitions and household dynamics in political participation (Stoker and Jennings, 1995; Dahlgaard, 2018; Dehdari et al., 2022).

4.1. Baseline specification

We begin by estimating a baseline linear probability model of the form:

$$voted_{it} = \alpha_i + \lambda_t + \beta\, has\_minor_{it} + X_{it}'\delta + \varepsilon_{it}, \qquad (1)$$

where X_{it} includes time-varying covariates: real (deflated) income, age and age squared, marital status, an indicator for residential mobility between elections, and the distance from the individuals residence to the assigned polling station. These controls proxy for time-varying life-cycle characteristics and logistical voting costs that may change as individuals age or relocate.

The coefficient β captures the average within-individual difference in turnout associated with having minor children, net of observable time-varying factors, abstracting from heterogeneity across parental ages and childrens life stages. We estimate this model by ordinary least squares and report standard errors clustered at both the individual and precinct levels using multi-way clustering. Given the binary outcome and the inclusion of individual fixed effects, the linear probability model provides directly interpretable marginal effects and avoids incidental parameters concerns common to nonlinear fixed-effects estimators.

4.2. Parental age and child life-cycle interactions

To assess whether the turnout gap between parents and non-parents varies with parental age, we estimate interaction models of the form:

$$voted_{it} = \alpha_i + \lambda_t + \gamma_1\, has\_minor_{it} + \gamma_2 \left(has\_minor_{it} \times age_{it}\right) + X_{it}'\delta + \varepsilon_{it}. \quad (2)$$

The interaction term tests whether the participation gap associated with parenthood evolves over the life cycle. To allow for non-linearities related to childrens development, we also estimate fully interacted models in which indicators for the youngest childs life stage are interacted with parental age. These specifications relax functional-form assumptions and allow the effects of parenthood to be concentrated at specific stages of childrens lives.

4.3. Sample restrictions and heterogeneity

The estimation sample mirrors the descriptive analysis. We restrict attention to individuals with valid turnout records and trim the age distribution by excluding observations below the first and above the 99th percentiles, computed among parents, to reduce the influence of extreme values. In heterogeneity analyses by gender, we estimate models separately for men and women. In additional specifications, we focus on women and compare mothers to childless women, with particular attention to mothers of children in early childhood (ages 0–5).

4.4. Identification and interpretation

Identification relies on within-individual variation over time. Conditional on individual fixed effects and election-year fixed effects, changes in parental status and in the age of children are assumed to be orthogonal to persistent unobserved determinants of voting propensity. This identification strategy follows a well-established approach in the literature on life-course transitions and household effects, where within-individual designs are used to document systematic participation gaps associated with major family events rather than to claim sharp causal effects.

The inclusion of time-varying controls mitigates concerns that family formation coincides mechanically with changes in income, residential mobility, or access to polling locations. Nonetheless, becoming a parent is a choice that may correlate with unobserved time-varying preferences or constraints. For this reason, the estimates are interpreted as descriptive within-individual associations that capture how participation systematically evolves with parental responsibilities over the life cycle, rather than as causal effects of parenthood.

We probe the robustness of the results by augmenting the control set, adding neighborhood fixed effects to absorb spatial heterogeneity in turnout, and estimating gender-specific and fully interacted models. Across these alternative specifications, the qualitative patterns are stable.

Finally, the empirical analysis focuses on participation outcomes. We do not directly estimate policy choices, political responsiveness, or government accountability. Accordingly, any implications for democratic representation and accountability are discussed as suggestive consequences of systematic participation gaps, rather than as estimated effects on policy outcomes.

5. Results

This section examines how parenthood affects electoral participation over the life cycle. We focus on within-individual estimates that allow the turnout gap associated with parenthood to vary with parental age and with the life stage of children. Unless otherwise stated, all specifications include individual and election-year fixed effects, and standard errors are clustered at the individual and precinct levels.

Conventional specifications that compare parents and non-parents without accounting for parental age yield little evidence of an average turnout gap. These benchmark results are reported in the Appendix for reference. We therefore concentrate on specifications that explicitly model life-cycle heterogeneity.

5.1. Parental age, electoral context, and child life stage

Table 2 presents the core results of the paper. The table progressively relaxes the assumption that the association between parenthood and turnout is homogeneous across the life cycle and across electoral contexts.

Columns 1 and 2 estimate models in which the effect of parenthood is allowed to vary with parental age. In both specifications, having a minor child is associated with a substantial reduction in turnout at younger ages. The positive and statistically significant interaction between parenthood and age indicates that this penalty

Table 2: Parenthood, age, and voter turnout: within-individual estimates

	(1) Voted	(2) Voted	(3) Voted	(4) Voted	(5) Voted
Has any minor child	-0.085***	-0.080***	-0.072***		
	(0.010)	(0.010)	(0.011)		
Age	0.124***	0.062***	0.062***	0.061***	0.062***
	(0.006)	(0.011)	(0.011)	(0.012)	(0.011)
Has any minor child × Age	0.002***	0.002***	0.002***		
	(0.000)	(0.000)	(0.000)		
Has any minor child × Municipal elections			-0.004**		
			(0.002)		
Has any child below age 5				-0.090***	
				(0.021)	
Has any child below age 5 × Age				0.002***	
				(0.001)	
Nursery					-0.050**
					(0.024)
Nursery × Age					0.001
					(0.001)
Preschool					-0.019
					(0.021)
Preschool × Age					0.000
					(0.001)
Elementary School					0.002
					(0.019)
Elementary School × Age					0.000
					(0.000)
Middle School					-0.023
					(0.022)
Middle School × Age					0.001
					(0.000)
High School					-0.055***
					(0.021)
High School × Age					0.001***
					(0.000)
Controls		✓	✓	✓	✓
Individual F.E.	✓	✓	✓	✓	✓
Year F.E.	✓	✓	✓	✓	✓
R^2	0.558	0.533	0.533	0.553	0.533
Observations	539,470	468,019	468,019	325,391	468,019

Notes: Linear probability models (OLS). The dependent variable is an indicator equal to one if the individual voted in the given election. All specifications include individual and election-year fixed effects. Column 1 estimates a baseline age-interacted parenthood model without additional controls. Column 2 adds time-varying controls for real income, age and age squared, residential mobility, distance to the assigned polling station, and marital status. Column 3 allows the parenthood–age gradient to differ between municipal and national elections by interacting parental status with a municipal-election indicator. Column 4 restricts the sample to parents of children below age five and interacts this indicator with parental age. Column 5 decomposes parenthood by the age of the youngest child and interacts each life-stage indicator with parental age. Standard errors are clustered at the individual and precinct levels. Coefficients capture within-individual associations and should not be interpreted as causal effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

attenuates steadily over the life cycle.

Evaluated at age 25, parents are approximately three percentage points less likely to vote than comparable non-parents. The implied gap narrows steadily with parental age and is close to zero by age 40. These estimates indicate that average comparisons between parents and non-parents obscure substantial life-cycle heterogeneity.

Column 3 addresses a concern regarding electoral context. It allows the parenthood–age gradient to differ between municipal and national elections by interacting parental status with an indicator for municipal contests. This specification tests whether the turnout penalty associated with early parenthood reflects differences in electoral salience, information requirements, or the perceived stakes of local versus national elections.

The estimates show that the age-dependent turnout penalty for parents is highly similar across electoral levels. While turnout is on average slightly lower in municipal elections, the additional interaction term is small in magnitude. This similarity across electoral levels suggests that the observed participation gap is unlikely to be driven by election-specific mobilization or differences in informational demands.

Column 4 restricts attention to parents of children below age five. The estimated turnout penalty at younger ages is considerably larger in magnitude and declines sharply with parental age. This pattern is consistent with the interpretation that the participation gap is driven by the intensity of childcare demands during early childhood.

Column 5 further decomposes parenthood by the age of the youngest child. The largest and most persistent participation gaps are observed among parents of infants and preschoolers. By contrast, coefficients for parents of school-age children are small and imprecisely estimated, and attenuate with parental age. These patterns reinforce the interpretation that intensive childcare responsibilities, rather than parenthood per se, drive the observed turnout gap.

5.2. Mothers and early childhood

The age-dependent participation gap documented above is not uniform across genders. Table 3 focuses on women and compares mothers to childless women, with particular attention to early childhood.

Once parental age is explicitly accounted for, the estimated turnout penalty for young mothers is sizable. At age 25, mothers of young children are substantially less likely to vote than comparable women without children. The positive interaction with age implies a rapid attenuation of this gap, which becomes negligible by around age 40. The magnitude of the participation penalty among young mothers is larger than in the full sample, indicating that women drive the aggregate results.

Table 3: Motherhood, age, and voter turnout

	(1) Voted	(2) Voted	(3) Voted
Has any child below age 5	-0.013***	-0.162***	-0.160***
	(0.005)	(0.034)	(0.034)
Age		0.073***	0.073***
		(0.020)	(0.020)
Has any child below age 5 × Age		0.004***	0.004***
		(0.001)	(0.001)
Controls	✓	✓	✓
Individual F.E.	✓	✓	✓
Year F.E.	✓	✓	✓
Neighborhood F.E.			✓
R^2	0.518	0.518	0.518
Observations	108,332	108,332	108,267

Notes: Linear probability models (OLS) estimated on women. The dependent variable is an indicator equal to one if the individual voted in the given election. The key regressor is an indicator for having a child below age five and its interaction with parental age. Controls include real income, age and age squared, residential mobility, distance to the assigned polling station, and marital status. Column 3 additionally includes distance to nurseries and preschools. Column 4 adds neighborhood fixed effects. All specifications include individual and election-year fixed effects. Standard errors are clustered at the individual and precinct levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

To aid interpretation of the interaction effects, Figure 2 plots the implied turnout gap between mothers and childless women across parental ages, holding other covariates at their sample means. The participation deficit is pronounced in early adulthood and declines smoothly with age.

The concentration of the turnout penalty among parents of 0–5 year-old, combined with the null effects of proximity to childcare facilities in robustness specifications, is most consistent with binding within-household time constraints rather than physical access to infrastructure or shifts in policy preferences per se.

Taken together, these results indicate that reduced electoral participation is concentrated among parents facing the most time-intensive stages of child-rearing, and is driven primarily by mothers. As parental age increases and childcare constraints ease, the participation gap dissipates.

6. Conclusion

This paper documents a pronounced life-cycle pattern in the association between parenthood and electoral participation. Using universe-level administrative data

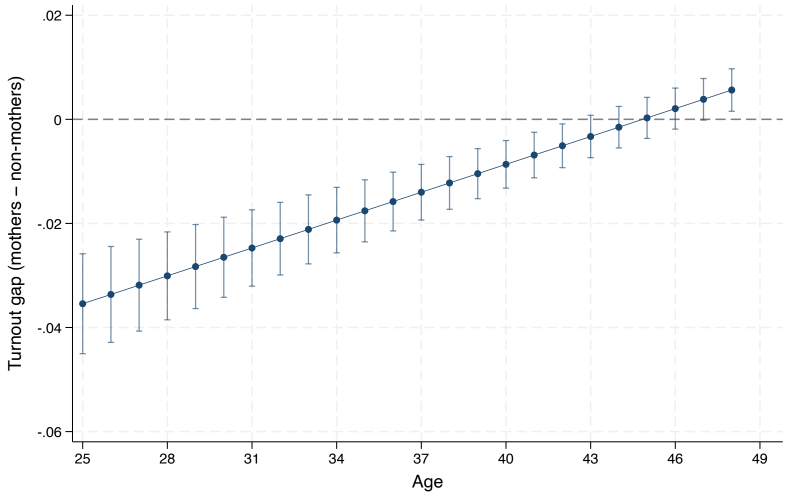

Fig. 2: Turnout gap between mothers and childless women by age

Notes: The figure plots predicted differences in turnout probabilities between women with children below age five and women without children, based on a linear probability model with individual and election-year fixed effects. Shaded areas represent 95% confidence intervals with standard errors clustered at the individual and precinct levels.

from the municipality of Bologna and exploiting within-individual variation over time, we show that parenthood is not associated with lower turnout on average. Instead, participation gaps are sharply concentrated among parents of very young children and decline steadily with parental age, disappearing by around age forty. The magnitude of these gaps—on the order of several percentage points at younger ages—declines smoothly over the life cycle.

Two features of the results are particularly salient. First, the participation penalty is concentrated in the most time-intensive phases of child-rearing, when children are infants or preschool-aged. Second, this pattern is driven almost entirely by women. Fathers turnout is largely unaffected by the presence or age of children, whereas young mothers exhibit substantial and age-dependent participation gaps. Together, these findings indicate that it is not parenthood per se, but the timing and intensity of childcare responsibilities, that shape electoral participation over the life course.

The empirical design does not identify causal effects of parenthood. Becoming a parent is a choice that may correlate with unobserved, time-varying preferences or constraints. Accordingly, the estimates are interpreted as descriptive within-individual associations rather than as causal effects. Nonetheless, the longitudinal structure of the data and the inclusion of individual fixed effects allow us to rule out explanations based on permanent differences in civic engagement between parents and non-parents. The stability of the results across specifications, subsamples, and electoral contexts further suggests that the documented patterns reflect systematic life-cycle regularities rather than artifacts of sample composition or residential sorting.

Although the analysis focuses on a single municipality, several features enhance the external relevance of the findings. Italy combines automatic voter registration, high baseline turnout, and mechanically assigned polling stations, limiting institutional barriers to participation. The observed turnout penalties therefore arise in a setting where formal voting costs are relatively low. If anything, this suggests that childcare-related time constraints may exert even stronger effects in contexts with higher voting costs or less supportive institutional environments. Moreover, the results closely align with evidence from other European settings, indicating that the mechanisms documented here are unlikely to be specific to Bologna.

6.1. Implications for democratic representation and accountability

The documented participation gaps among young parents—particularly mothers of infants and preschool-aged children—have potentially important implications for democratic representation and political accountability. Periods of reduced participation imply that individuals with distinct preferences over public goods such as childcare, early education, housing, and family policies may be systematically underrepresented precisely when family-related constraints are most binding.

A large literature emphasizes that political participation shapes both electoral accountability and the allocation of public resources. When turnout is lower among specific demographic groups, politicians face weaker incentives to cater to their preferences, even in the absence of formal disenfranchisement (Baland and Robinson, 2003; Fujiwara, 2015). In this context, lower participation among young parents may translate into three related implications.

First, political accountability may be attenuated with respect to family-oriented services. If young parents—and especially mothers—are less likely to vote during periods of intensive childcare, elected officials may have weaker incentives to expand or improve family-friendly policies such as flexible childcare slots, extended opening hours, or early-childhood services. Second, the composition of public spending may tilt toward policy areas favored by more politically active groups, such as older voters, potentially biasing budgets toward health care and pensions at the expense of early-childhood investments. Third, reduced participation among young mothers may contribute to persistent gender imbalances in political representation, including lower female candidacy and weaker descriptive representation in policy areas directly related to family welfare.

While the administrative data used in this paper do not permit direct testing of these channels, the observed patterns are consistent with evidence from other institutional settings. Existing studies document that groups with lower electoral participation exert less influence on policy outcomes and public spending priorities (Baland and Robinson, 2003; Fujiwara, 2015). In low-fertility societies where the age at first birth has risen substantially, the temporary participation gaps documented here may therefore have non-trivial aggregate consequences as the window of reduced parental participation shifts later in the life cycle.

Importantly, these implications should be interpreted as suggestive rather than causal. Nonetheless, they highlight a mechanism through which demographic change and family dynamics may interact with political participation to shape democratic representation and accountability, even in institutional environments with low formal barriers to voting.

More broadly, this study highlights the value of linking administrative electoral records to household and demographic data. By allowing researchers to follow the same individuals across major life transitions, such data make it possible to document participation patterns that are difficult to detect using cross-sectional surveys. In the context studied here, this approach reveals that participation gaps associated with parenthood are not permanent, but concentrated in specific phases of the life cycle. Future research could extend this framework to other institutional settings, examine cohort-specific dynamics, or explore how temporary participation gaps translate into longer-term political inequalities.

References

[1] Akee, R., Copeland, W., Costello, E. J., Holbein, J. B., and Simeonova, E. (2018). Family income and the intergenerational transmission of voting behavior: Evidence from an income intervention. *NBER Working Paper*, (24770).

[2] Ansolabehere, S. and Hersh, E. (2012). Validation: What big data reveal about survey misreporting and the real electorate. *Political Analysis*, 20(4):437–459.

[3] Bailey, A. K. (2009). How personal is the political? democratic revolution and fertility decline. *Journal of Family History*, 34(4):409–428.

[4] Baland, J.-M. and Robinson, J. A. (2003). Oligarchy, democracy and property rights. *Journal of Public Economics*, 87(7):1399–1425.

[5] Bellettini, G., Berti Ceroni, C., Cantoni, E., Monfardini, C., and Schafer, J. (2023). Modern family? the gendered effects of marriage and childbearing on voter turnout. *British Journal of Political Science*, 53(3):10161040.

[6] Bhatti, Y., Hansen, K. M., Naurin, E., Stolle, D., and Wass, H. M. (2019). Can you deliver a baby and vote? the effect of the first stages of parenthood on voter turnout. *Journal of Elections, Public Opinion and Parties*, 29(1):61–81.

[7] Boix, C., Magyar, Z., and Muñoz, J. (2025). Turnout, family, and gender norms:

The political incorporation of women in sweden, 1921–1960. *World Politics*, 77(2):242–287.

[8] Brulé, R. and Gaikwad, N. (2021). Culture, capital, and the political economy gender gap: Evidence from meghalayas matrilineal tribes. *Journal of Politics*, 83(3):834–850.

[9] Cantoni, E. (2020). A precinct too far: Turnout and voting costs. *American Economic Journal: Applied Economics*, 12(1):61–85.

[10] Cantoni, E., Gazzè, L., and Schafer, J. (2021). Turnout in concurrent elections: Evidence from two quasi-experiments in italy. *European Journal of Political Economy*, 70:102035.

[11] Cheema, A., Khan, S., Liaqat, A., and Mohmand, S. K. (2023). Canvassing the gatekeepers: A field experiment to increase women voters turnout in pakistan. *American Political Science Review*, 117(1):121.

[12] Cutts, D. and Fieldhouse, E. (2009). What small spatial scales are relevant as electoral contexts for individual voters? the importance of the household on turnout at the 2001 general election. *American Journal of Political Science*, 53(3):726–739.

[13] Dahlgaard, J. O. (2018). Trickle-up political socialization: The causal effect on turnout of parenting a newly enfranchised voter. *American Political Science Review*, 112(3):698–705.

[14] Dahlgaard, J. O., Bhatti, Y., Hansen, J. H., and Hansen, K. M. (2022). Living together, voting together: Voters moving in together before an election have higher turnout. *British Journal of Political Science*, 52(2):631–648.

[15] Dassonneville, R. and Kostelka, F. (2021). The cultural sources of the gender gap in voter turnout. *British Journal of Political Science*, 51(3):1040–1061.

[16] Dehdari, S. H., Lindgren, K., Oskarsson, S., and Vernby, K. (2022). The ex-factor: Examining the gendered effect of divorce on voter turnout. *American Political Science Review*, 116(4):1293–1308.

[17] Downs, A. (1957). *An Economic Theory of Democracy.* Harper and Row, New York.

[18] Fujiwara, T. (2015). Voting technology, political responsiveness, and infant health: Evidence from brazil. *Econometrica*, 83(2):423–464.

[19] Grechyna, D. (2023). Parenthood and political engagement. *European Journal of Political Economy*, 76:102238.

[20] Gruneau, M. F. (2018). Reconsidering the partner effect on voting. *Electoral Studies*, 53:48–56.

[21] Gruneau, M. F. (2020). Assortative mating and turnout: A self-reinforcing pattern of unequal voting participation. *European Political Science Review*, 12(2):155–171.

[22] Inglehart, R. and Norris, P. (2003). *Rising Tide: Gender Equality and Cultural Change Around the World*. Cambridge University Press.

[23] Iversen, T. and Rosenbluth, F. (2006). The political economy of gender: Explaining cross-national variation in the gender division of labour and the gender voting gap. *American Journal of Political Science*, 50(1):1–19.

[24] Jennings, M. K. (1983). Gender roles and inequalities in political participation: Results from an eight-nation study. *Western Political Quarterly*, 36(3):364–385.

[25] Kittilson, M. C. (2016). Gender and political behavior. In *Oxford Research Encyclopedia of Politics*. Oxford University Press. Research encyclopedia article.

[26] Milazzo, A. and Goldstein, M. (2019). Governance and womens economic and political participation: Power inequalities, formal constraints and norms. *The World Bank Research Observer*, 34(1):34–64.

[27] Miller, G. (2008). Womens suffrage, political responsiveness, and child survival in american history. *Quarterly Journal of Economics*, 123(3):1287–1327.

[28] Prillaman, S. A. (2023). Strength in numbers: How womens groups close indias political gender gap. *American Journal of Political Science*, 67(2):390–410.

[29] Quaranta, M. and Dotti Sani, G. M. (2018). Left behind? gender gaps in political engagement over the life course in twenty-seven european countries. *Social Politics: International Studies in Gender, State & Society*, 25(2):254–286.

[30] Schafer, J., Cantoni, E., Bellettini, G., and Berti Ceroni, C. (2022). Making unequal democracy work? the effects of income on voter turnout in northern italy. *American Journal of Political Science*, 66(3):745–761.

[31] Schlozman, K. L., Burns, N., and Verba, S. (1994). Gender and the pathways to participation: The role of resources. *Journal of Politics*, 56(4):963–990.

[32] Smets, K. and van Ham, C. (2013). The embarrassment of riches? a meta-analysis of individual-level research on voter turnout. *Electoral Studies*, 32(2):344–359.

[33] Sommer, U. (2018). Women, demography, and politics: How lower fertility rates lead to democracy. *Demography*, 55(2):559–583.

[34] Stoker, L. and Jennings, M. K. (1995). Life-cycle transitions and political participation: The case of marriage. *American Political Science Review*, 89(2):421–433.

[35] Teele, D. L. (2018). *Forging the Franchise: The Political Origins of the Womens Vote*. Princeton University Press.

[36] Teney, C., Dochow-Sondershaus, S., and Lovette, F. (2024). The gendered effect of parenthood on voting behaviour in the 2021 german federal election. *German Politics*, 33(1):22–45.

[37] Verba, S., Burns, N., and Schlozman, K. L. (1997). Knowing and caring about politics: Gender and political engagement. *Journal of Politics*, 59(4):1051–1072.

[38] Voorpostel, M. and Coffé, H. (2012). Transitions in partnership and parental status, gender, and political and civic participation. *European Sociological Review*, 28(1):28–42.

[39] Wolfinger, R. E. and Rosenstone, S. J. (1980). *Who Votes?* Yale University Press.

Appendix

A. Related Literature and Institutional Context

This appendix situates the paper within the broader literature on political participation, gender, family dynamics, and democratic representation. The discussion complements the main text and provides additional context for the empirical analysis.

Gender norms, institutions, and political participation. A large literature documents persistent gender gaps in political engagement and emphasizes the role of social norms, labor-market institutions, and caregiving responsibilities (Jennings, 1983; Schlozman et al., 1994; Verba et al., 1997). Cross-national evidence links variation in gender gaps to welfare-state institutions and labor-market structures (Iversen and Rosenbluth, 2006; Kittilson, 2016; Dassonneville and Kostelka, 2021). Long-run historical analyses show how womens enfranchisement and political incorporation affected policy responsiveness and social outcomes (Miller, 2008; Teele, 2018), while recent work highlights the persistence of gender norms even in advanced democracies (Inglehart and Norris, 2003; Boix et al., 2025). Evidence from developing contexts further illustrates how institutional constraints and patriarchal norms limit womens civic engagement (Milazzo and Goldstein, 2019).

Household dynamics and life-course transitions. Political participation is shaped by household structure and family transitions. Marriage increases turnout through socialization and mobilization within households (Stoker and Jennings, 1995), while partnership formation and dissolution affect participation in gendered ways (Voorpostel and Coffé, 2012; Dehdari et al., 2022). Living arrangements matter: cohabitation with politically engaged partners and moving in together before elections raise turnout (Cutts and Fieldhouse, 2009; Dahlgaard et al., 2022). Assortative mating can reinforce inequalities in participation (Gruneau, 2018, 2020). Intergenerational mechanisms also operate in reverse: parenting newly enfranchised voters can increase parents turnout (Dahlgaard, 2018). Economic resources and shocks further shape participation (Akee et al., 2018; Schafer et al., 2022).

Parenthood, childcare, and comparative evidence. Recent studies analyze how parenthood affects political participation. Survey-based and administrative evidence documents a decline in womens turnout following childbirth, with smaller or null effects for men (Quaranta and Dotti Sani, 2018; Bhatti et al., 2019; Grechyna, 2023; Teney et al., 2024). These effects are strongest in early childhood and attenuate over time, consistent with childcare-related opportunity costs. Comparative research highlights substantial institutional heterogeneity: mobilization strategies targeting household gatekeepers can increase womens turnout in patriarchal contexts (Cheema et al., 2023), while collective organizations strengthen womens civic skills and participation (Prillaman, 2023). Cultural institutions can even reverse gender gaps, as shown in matrilineal societies (Brulé and Gaikwad, 2021).

Overall, this literature emphasizes the role of gendered opportunity costs and household dynamics in shaping political participation. By exploiting universe-level administrative data and distinguishing parental age from childrens life stages, the present study complements this work.

B. Descriptive Statistics and Additional Figures

This appendix reports additional descriptive statistics and figures that complement the main analysis. All statistics are based on raw means and are intended to document aggregate turnout patterns rather than to identify causal effects. The figures illustrate how voter participation varies across age, parental status, and child life stages, and serve to motivate the within-individual specifications used in the main text.

B.1. Descriptive statistics

Table B.1, Table B.2, and Table B.3 report summary statistics for the full sample and for selected subsamples. Parents differ systematically from non-parents along observable dimensions such as age and marital status, while income distributions and distances to polling stations are broadly comparable. These differences underscore the importance of controlling for life-cycle factors and individual fixed effects in the regression analysis.

Full sample. Table B.1 shows how parents are younger and substantially more likely to be married than non-parents, while income distributions are broadly comparable

Table B.1: Descriptive statistics — full sample

	(1) No children					(2) With children				
	mean	sd	min	max	count	mean	sd	min	max	count
Female	0.536	0.499	0.000	1.000	1106962	0.555	0.497	0.000	1.000	198000
Age	54.889	20.002	18.000	113.000	1106962	42.307	7.137	18.000	65.000	198000
Real Income	23490.394	36886.829	0.000	8616915.000	938380	29322.967	44121.146	0.000	6039291.000	173393
Distance	4.771	4.061	0.002	77.727	1106030	5.055	4.551	0.002	75.330	197768
moved	0.064	0.245	0.000	1.000	1106962	0.110	0.313	0.000	1.000	198000
married	0.451	0.498	0.000	1.000	1106962	0.867	0.340	0.000	1.000	198000
Municipal election (vs national)	0.500	0.500	0.000	1.000	1106962	0.494	0.500	0.000	1.000	198000

Notes: Descriptive statistics for adults aged 18–65. "With children" indicates the presence of at least one cohabiting minor. Income is deflated to constant euros. Distances are measured in kilometers from the individuals residence to the assigned polling station.

across groups. Differences in distance to polling stations are small, indicating similar geographic access to voting locations.

Female sample. In Table B.2, among women, mothers and non-mothers display similar income levels, while mothers are more likely to be married and slightly more likely to have moved between elections.

Table B.2: Descriptive statistics — female sample

	(1) No children					(2) With children				
	mean	sd	min	max	count	mean	sd	min	max	count
Age	40.921	9.538	25.000	56.000	235851	41.138	6.599	25.000	56.000	107969
Real Income	20583.429	22302.937	0.000	2137710.250	190109	21317.620	22914.840	0.000	1220693.750	89592
Distance	4.807	4.128	0.002	77.727	235612	5.041	4.549	0.002	75.330	107845
moved	0.101	0.301	0.000	1.000	235851	0.111	0.314	0.000	1.000	107969
married	0.362	0.481	0.000	1.000	235851	0.832	0.374	0.000	1.000	107969
Municipal election (vs national)	0.500	0.500	0.000	1.000	235851	0.492	0.500	0.000	1.000	107969

Notes: Descriptive statistics for women aged 25–55. Variable definitions follow those in the main text.

Female sample with young children. The subsample showed in Table B.3 isolates women in the most time-intensive phases of child-rearing and illustrates the sharper life-cycle differences in age, marital status, and residential mobility that motivate the early-childhood specifications in the main analysis.

Table B.3: Descriptive statistics — female sample with children at nursery and preschool

	(1) No children					(2) With children				
	mean	sd	min	max	count	mean	sd	min	max	count
Age	34.904	7.062	23.000	48.000	169389	36.270	5.125	23.000	48.000	43090
Real Income	18275.680	17700.362	0.000	954396.562	132101	18715.793	22207.242	0.000	1220693.750	34191
Distance	4.775	4.120	0.002	77.727	169192	5.025	4.506	0.002	70.584	43021
moved	0.114	0.318	0.000	1.000	169389	0.150	0.357	0.000	1.000	43090
married	0.228	0.419	0.000	1.000	169389	0.819	0.385	0.000	1.000	43090
Municipal election (vs national)	0.500	0.500	0.000	1.000	169389	0.500	0.500	0.000	1.000	43090

Notes: Descriptive statistics for women aged 25–48, comparing mothers of children aged 0–5 to childless women.

B.2. Additional figures

The figures in this subsection plot raw turnout averages by age and parental status for different subsamples. They are purely descriptive and do not exploit the panel structure of the data. Accordingly, they should not be interpreted as evidence of causal effects of parenthood on turnout. Their purpose is to visualize life-cycle patterns and to illustrate the heterogeneity that motivates the interaction models estimated in the main text.

Age quintiles. Figure B.1 plots average turnout by age quintile for individuals with and without cohabiting minor children. Turnout is high across all age groups. Parents exhibit lower participation in the youngest age quintile, while differences

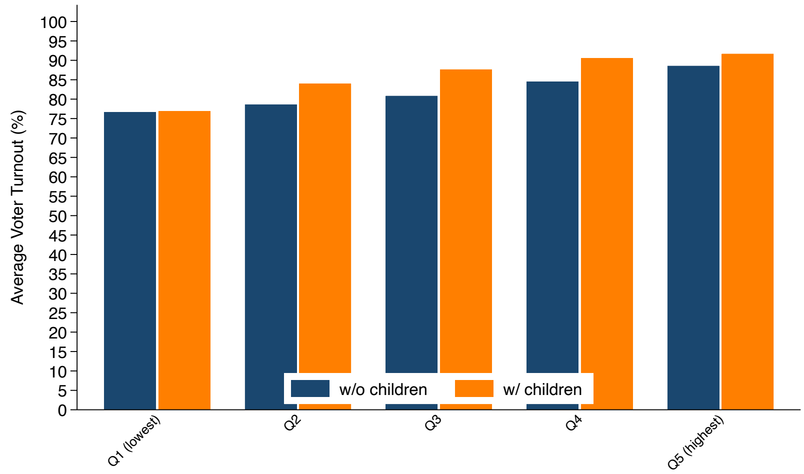

Fig. B.1: Average voter turnout by age quintile and parental status

Notes: Raw turnout means by age quintiles for individuals with and without cohabiting minor children. Based on the restricted sample (ages 25–56).

narrow substantially at older ages. These patterns are consistent with the life-cycle patterns documented in the regression analysis.

Women. Figure B.2 and Figure B.3 replicate the age–turnout profiles for the female subsample. Turnout increases with age for both mothers and non-mothers. At younger ages, mothers display lower participation, while the gap narrows progressively over the life cycle. These descriptive patterns are consistent with the age-dependent participation gaps estimated in the main results.

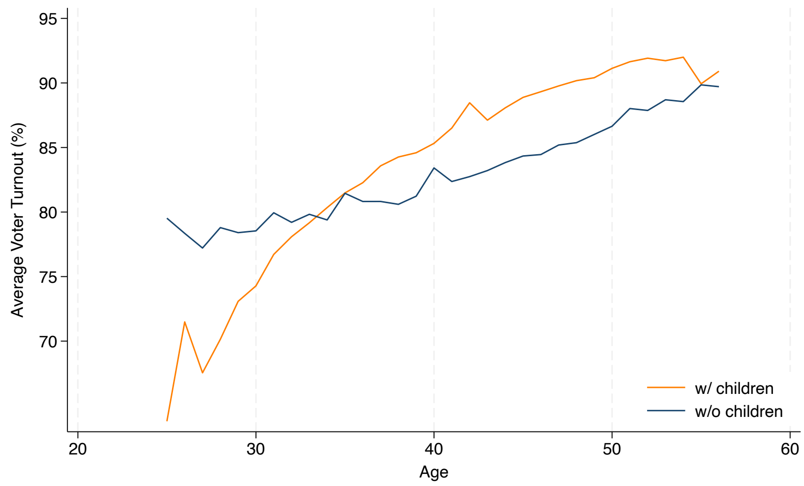

Fig. B.2: Average voter turnout by age and parental status (female sample)

Notes: Raw turnout means by age for women with and without cohabiting minor children. Based on the restricted female sample (ages 25–55).

Women with young children. Figure B.4 and Figure B.5 focus on women with at least one child aged 0–5. Across all ages, mothers of young children exhibit

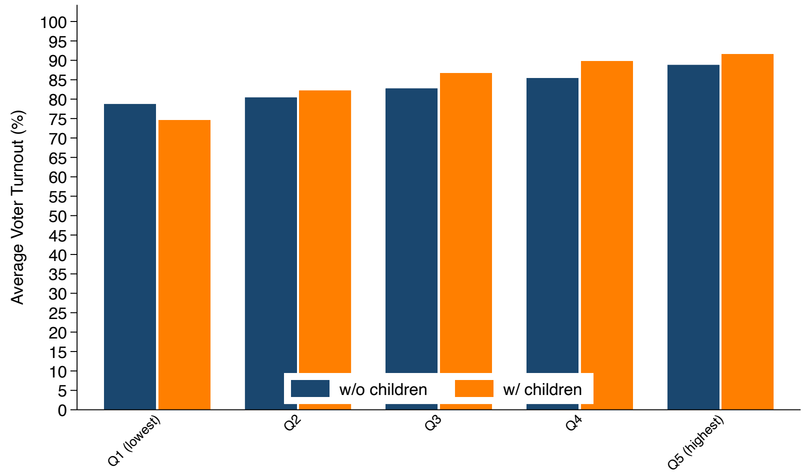

Fig. B.3: Average voter turnout by age quintile and parental status (female sample)

Notes: Raw turnout means by age quintiles for women with and without cohabiting minor children. Based on the restricted female sample (ages 25–55).

systematically lower turnout than childless women. The gap is largest at younger ages and declines thereafter. This visualization highlights the concentration of participation gaps during the most time-intensive phases of child-rearing.

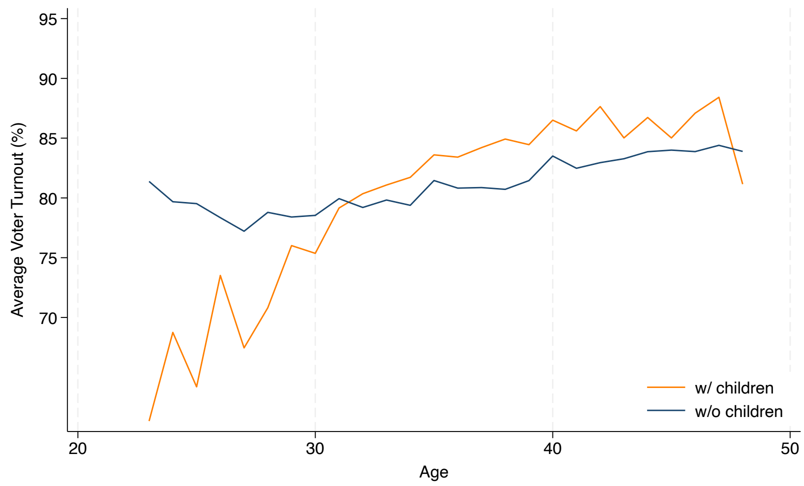

Fig. B.4: Average voter turnout by age (female sample, nursery and preschool)

Notes: Raw turnout means by age for women with at least one child aged 0–5 and women without children. Based on the restricted subsample (ages 25–48).

Taken together, these descriptive figures illustrate the raw turnout patterns that motivate the within-individual analysis. For formal inference, the analysis therefore relies on the within-individual estimates reported in the main text.

C. Baseline and Robustness Checks

This appendix reports benchmark specifications and robustness checks that complement the main analysis. All models are estimated as linear probability models

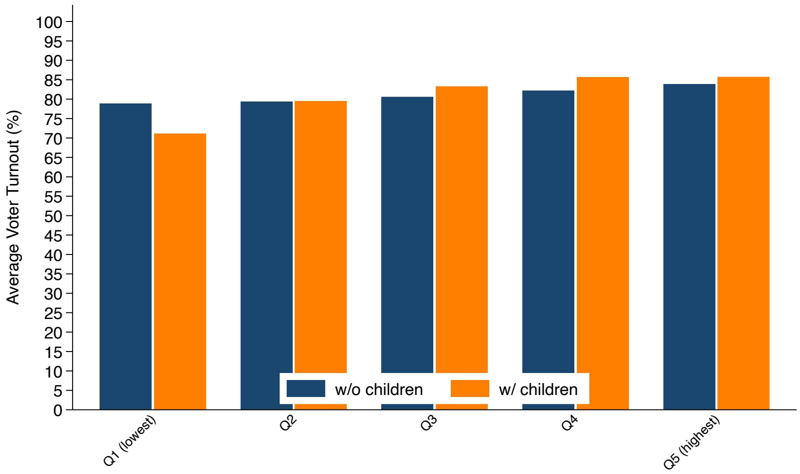

Fig. B.5: Average voter turnout by age quintile (female sample, nursery and preschool)

Notes: Raw turnout means by age quintiles for women with at least one child aged 0–5 and women without children. Based on the restricted subsample (ages 25–48).

(OLS) with individual and election-year fixed effects, unless otherwise stated. Standard errors are clustered at the individual and precinct levels.

The purpose of this appendix is twofold. First, it documents that conventional benchmark specifications that do not account for life-cycle heterogeneity yield little evidence of an average turnout gap between parents and non-parents. Second, it shows that the age-dependent participation patterns documented in the main text are robust to alternative samples and specifications.

C.1. Benchmark estimates

Table C.1 reports benchmark estimates that compare parents and non-parents without allowing the turnout gap to vary with parental age.

Consistent with previous studies, Column 1 shows no statistically significant average difference in turnout between parents and non-parents once individual fixed effects are included. Column 2 suggests lower participation among parents of very young children and slightly higher turnout among parents of older children. These patterns reflect the tight correlation between parental age and childrens age and motivate the age-interaction specifications in the main text.

C.2. Robustness checks

Table C.2 assesses whether the main results are sensitive to additional controls and spatial heterogeneity.

The estimated turnout penalty at younger ages and its attenuation with parental age remain quantitatively and statistically similar across specifications. These re-

Table C.1: Benchmark estimates: parenthood and voter turnout

	(1) Voted	(2) Voted
Has any minor child	-0.002	
	(0.002)	
Age	0.064***	0.062***
	(0.011)	(0.011)
Nursery		-0.018***
		(0.003)
Preschool		-0.004*
		(0.002)
Elementary School		0.004*
		(0.002)
Middle School		0.005**
		(0.002)
High School		0.006***
		(0.002)
Controls	✓	✓
Individual F.E.	✓	✓
Year F.E.	✓	✓
R^2	0.533	0.533
Observations	468,019	468,019

Notes: Linear probability models (OLS) with individual and election-year fixed effects. The dependent variable is an indicator equal to one if the individual voted in the given election. Column 1 includes an indicator for having any minor child. Column 2 distinguishes parents by the age of the youngest child. Controls include real income, age and age squared, residential mobility, distance to the polling station, and marital status. Standard errors are clustered at the individual and precinct levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.2: Robustness checks

	(1) Voted	(2) Voted
Has any minor child	-0.080***	-0.081***
	(0.010)	(0.010)
Age	0.062***	0.063***
	(0.011)	(0.011)
Has any minor child × Age	0.002***	0.002***
	(0.000)	(0.000)
Controls	✓	✓
Individual F.E.	✓	✓
Year F.E.	✓	✓
Neighborhood F.E.	✓	
Additional Control: Population density		✓
R^2	0.533	0.533
Observations	467,724	466,768

Notes: Linear probability models (OLS) with individual and election-year fixed effects. Column 1 adds neighborhood fixed effects, absorbing time-invariant spatial heterogeneity in turnout. Column 2 additionally controls for neighborhood population density. Baseline controls follow the main text. Standard errors are clustered at the individual and precinct levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

sults indicate that the main findings are not driven by residential sorting, neighborhood composition, or omitted spatial factors.

C.3. Heterogeneity by gender and marital status

Table C.3 examines whether the age-dependent turnout penalty varies systematically by gender and marital status.

The participation gap associated with parenthood is absent for men and present for women, regardless of marital status. While the magnitude of the estimates differs across subsamples, the qualitative pattern is unchanged. These results mitigate concerns that the main findings are driven by compositional differences between parents and non-parents or by marital status.

Table C.3: Robustness by gender and marital status

	Men Subsample	Women Subsample	Non-Married Subsample	Married Subsample
	(1) Voted	(2) Voted	(3) Voted	(4) Voted
Has any minor child	-0.027*	-0.121***	-0.074**	-0.069***
	(0.014)	(0.013)	(0.029)	(0.012)
Age	0.052***	0.074***	0.046***	0.085***
	(0.015)	(0.016)	(0.015)	(0.019)
Has any minor child × Age	0.001**	0.003***	0.002***	0.001***
	(0.000)	(0.000)	(0.001)	(0.000)
Controls	✓	✓	✓	✓
Individual F.E.	✓	✓	✓	✓
Year F.E.	✓	✓	✓	✓
R^2	0.561	0.502	0.557	0.518
Observations	232,848	235,171	202,985	247,904

Notes: Linear probability models (OLS) with individual and election-year fixed effects, estimated separately by gender and marital status. Baseline controls follow the main text. Standard errors are clustered at the individual and precinct levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.