

LETTER

# Evaluating the Minority Candidate Penalty with a Regression Discontinuity Approach

Ariel White<sup>1</sup> , Paru Shah<sup>2</sup> , Eric Gonzalez Juenke<sup>3</sup> and Bernard L. Fraga<sup>4</sup> 

<sup>1</sup>Department of Political Science, Massachusetts Institute of Technology, Cambridge, MA, USA, <sup>2</sup>Department of Political Science, Rutgers University, New Brunswick, NJ, USA, <sup>3</sup>Department of Political Science, Michigan State University, East Lansing, MI, USA and <sup>4</sup>Department of Political Science, Emory University, Atlanta, GA, USA

**Corresponding author:** Paru Shah; Email: [paru.shah@rutgers.edu](mailto:paru.shah@rutgers.edu)

(Received 4 October 2022; revised 4 July 2023; accepted 10 September 2023; first published online 8 January 2024)

## Abstract

Do parties face an electoral penalty when they nominate candidates of colour? We employ a regression discontinuity design using state legislative election data from 2018, 2019, and 2020 to isolate the effect of nominating a candidate of colour on a party's general election performance. Utilising this approach with real-world data heightens external validity relative to existing racial penalty studies, largely supported by surveys and experiments. We find no evidence that candidates of colour are disadvantaged in state legislative general elections relative to narrowly nominated white candidates from the same party. These findings challenge the leading explanations for the underrepresentation of racial/ethnic minority groups, with implications for candidate selection across the United States.

**Keywords:** race/ethnic representation; candidate evaluations; voter bias

## Racial Bias and Election Outcomes

Many political scholars and practitioners believe that parties pay a significant penalty in the voting booth when they nominate racial and ethnic minority candidates, particularly in places with a heavily-white electorate (Kinder and Dale-Riddle 2012; Lewis-Beck, Tien, and Nadeau 2010). There are good reasons to believe this. Decades of research into white racism suggest that white voters are prejudiced against candidates of colour in experimental and survey settings. This prejudice results in a dearth of minority representation in majority-white electoral districts. This research is incredibly consequential, influencing everything from redistricting to candidate ambition to party leaders tasked with recruiting individuals to run for office (Canon and Posner 1999; Doherty, Dowling, and Miller 2019).

However, these electoral penalties may be overstated, especially in an era of partisan and racial polarization (Kam 2007). While candidates of colour may lose votes from some out-partisan voters, they may simultaneously get a bigger boost from some co-partisans, either in turnout or vote choice (Agadjanian et al. 2020; Chudy 2021; Chudy, Piston, and Shipper 2019; Jackman and Vavreck 2010; Stout 2020; Tesler and Sears 2010). Despite this recent evidence of potential electoral *advantages* to racial and ethnic minority candidates, scholars have yet to build a solid empirical record to evaluate its merit in the context of the racial penalty literature. In this research note, we offer new evidence to help adjudicate between these contrasting narratives about racial penalties.

Early research into white racial bias and elections in the United States focused on the paucity of minority legislators outside of majority-minority districts (Canon and Posner 1999; Lublin

© The Author(s), 2024. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

1999). This work is foundational in redistricting disputes and has been codified into representational legal frameworks for decades (Page v. Bartels 2001; Georgia v. Ashcroft 2003; Bartlett v. Strickland 2009). However, selection bias limits our ability to draw useful inferences from research that only looks at election winners (Juenke 2014), while work that includes both candidates and officeholders finds little evidence of bias that penalizes Black or Latino general election candidates that is distinct from the party (Fraga, Shah, and Juenke 2020; Juenke and Shah 2016; Voss and Lublin 2001). Furthermore, as candidates of colour do not uniformly emerge in all types of districts, it is empirically difficult to distinguish candidate supply versus voter demand effects using observational designs (Branton 2009) or to separate the effects of party from district racial and ethnic factors (Juenke and Shah 2016). Observational research struggles to directly test causal relationships between voters' attitudes and their ballot choices.

Given the limitations of observational work, some scholars have turned to experimental research that tests for racial bias in a controlled setting. Experimental work shows racial and ethnic considerations heavily influence white voters' candidate preferences. For example, information search processes differ for many white voters, harming the electoral prospects of candidates of colour (Crowder-Meyer, Gadarian, and Trounstein 2020a; Crowder-Meyer et al. 2020b; Ditonto 2020; Valentino, Hutchings, and White 2002). Partisanship can also work to either minimize the impact of racism (Kam 2007; Sigelman et al. 1995) or reinforce prejudice (Stout 2020; Tesler 2013). Similar disagreements abound regarding the interaction of candidate race and ethnicity and information like ideology, candidate quality, or incumbency (Andersen and Junn 2010; Kam 2007). The difficulty of replicating the complex and dynamic information environment of political campaigns limits the external validity of the evidence from these experimental studies. Thus, there remain important disagreements as to why we find so many bigots in the voting public and yet few aggregate effects across elections in various partisan and racial/ethnic contexts (Citrin, Green, and Sears 1990; Fraga, Shah, and Juenke 2020; Highton 2004; Voss and Lublin 2001).

This paper uses a regression discontinuity (RD) design that leverages closely fought primary elections yielding either white or non-white party nominees. We follow previous work using close elections to examine otherwise similar districts where different candidates are chosen (Broockman 2014; Fraga and Hassell 2021). RD designs emphasize real-world external validity while improving our ability to test causal claims. They are particularly valuable tools to help assess the disparate methodological evidence in this literature.

Like Bucchianeri's (2018) work on candidate gender, we leverage close primary winners to estimate the causal effect of party primary voters nominating either a white candidate or a candidate of colour. Because these nominees emerge from similar districts and primary elections near the cut point, we can more precisely and confidently estimate any differences in the subsequent two-party vote share in the subsequent general election. While we take care to consider the limitations of this design (De la Cuesta and Imai 2016; Marshall 2021), the results offer a unique and significant contribution to the racial penalty literature.

The novel data and design we use allow for a strong test of the racial penalty narrative and lead us to conclude that existing analyses provide an incomplete picture of electoral reality today. Indeed, our evidence supports more recent theories of racial polarization that suggest candidates of colour from either party may not face general election penalties in the voting booth precisely because of the extreme nature of racial polarization in the American public (Agadjanian et al. 2020; Chudy 2021; Chudy, Piston, and Shipper 2019; Jackman and Vavreck 2010; Stout 2015; Tesler and Sears 2010). The findings suggest several new directions for resolving the continued puzzle of minority success in a racially polarized electoral environment. The results may also help address a larger set of questions around increasing descriptive representation and more polarized backsliding on substantive representation in the US in the twenty-first century (Enders and Scott 2019; Grumbach 2022), which we consider in the conclusion.

## States as Laboratories of (Racial) Representation

State legislative elections are ideal for testing theories of bias because there are thousands of partisan elections in numerous states every two years compared to hundreds of congressional races, thus creating statistical power for the kinds of racial/ethnic comparisons that are difficult to make using congressional election data. This paper focuses on primary elections where a white candidate narrowly defeated or was narrowly defeated by a candidate of colour. The dataset consists of 400 elections (from forty-two states) in which the top two candidates in a primary were of different race or ethnicity, such that the outcome of the primary election determined whether the party's nominee was white or a candidate of colour.<sup>1</sup>

The regression discontinuity (RD) design provides a test of what happens to a party's general election vote share when a candidate of colour [rather than a white candidate] wins a competitive primary, and the party advances them to a competitive general election, relative to a white candidate [candidate of colour]. While other work has explored what happens when candidates of colour (Juenke and Shah 2016) and women of colour (Fraga, Shah, and Juenke 2020) emerge more generally, our analysis goes one step further in trying to parse out the general election *effect* of nominating candidates of colour after close primary elections. In sum, we use real-world data to examine the local average treatment effect of a barely winning (in the primary) white candidate or a candidate of colour on parties' general election outcomes.

We focus on general elections where the top two vote-getters in the preceding primary election were of different racial/ethnic backgrounds. That is, the candidate who received the party's nomination via a primary election was of a different race/ethnicity than the candidate who came in second place. The forcing variable in this design is the 'minority candidate primary win margin': what share did the minority candidate win (or lose) by? For example, if there were two candidates and the candidate of colour won 55 per cent of the vote and the white candidate 45 per cent, this variable would take the value 0.1. In practice, most instances where there is a difference in the race/ethnicity of the top two vote-getters occur when one of the candidates is [non-Hispanic] white. This aligns with previous literature and our substantive interest in examining the impact of nominating a candidate of colour relative to a white candidate.<sup>2</sup>

The dependent variable is the two-party general election vote share for the party that contested the interracial primary. Thus, if the Democrats were the party with the contested interracial primary, our dependent variable is the two-party Democratic vote share in the subsequent general election. In the case of a contested interracial Republican primary, the dependent variable is the inverse (that is, the Republican share of the two-party vote). For uncontested general elections, the general election two-party vote share is not meaningful. Thus, these observations are excluded from the analysis (see the Online Appendix for analyses that impute vote counts for these races).

## Results

We begin by graphically displaying the comparison being made in the RD design. [Figure 1](#) presents a binned scatterplot of the data described above.<sup>3</sup> Observations to the left of the vertical line marking the cut-off are primary elections where the minority primary win margin is negative; that is, the minority candidate lost the primary election and did not proceed to the general election. To the right of the cut-off line are cases in which the minority candidate won the primary and became their party's nominee in the election. The y-axis displays the mean values of our outcome measure,

<sup>1</sup>These races are drawn from a broader dataset developed by the Candidate Characteristics Collaborative (C3), a group of scholars who cooperatively coded the race, ethnicity, and gender of state legislative candidates in 2018 and 2020 (Fraga, Juenke, and Shah 2021), combined with the 2019 data collected by the authors. The dataset provides racial and ethnic information on thousands of candidates for office.

<sup>2</sup>See the Online Appendix for additional details about our election data, including a discussion of the partisan breakdown of matchups, how we dealt with multi-member and non-partisan primaries, and the exclusion of uncontested elections.

<sup>3</sup>Figure and all RD analyses produced with R package `rdrubust` (Calonico *et al.* 2021).




Figure 1. Illustration of the RDD setup showing the full range of the dataset.

the party's general election vote share for each bin. The curving line shows a polynomial function fit to the data on each side of the cut point. The larger black points are binned means, and, the light grey dots plot each individual observation (each election) in the dataset.

The apparent discontinuity at the cut point indicates that when a party narrowly nominates a non-white candidate for the general election instead of a white candidate, the party wins a slightly *larger* vote share in the general election that follows (not a smaller share as racial-penalty accounts would indicate). The intuition of the RD design is that observations proximate to the cut point should be very similar on all other observable and unobservable dimensions, and district-level characteristics are smooth across the cut point with no apparent discontinuities (as shown in Figure A4 of the Online Appendix). The main difference is that the party's voters either do or do not nominate the white candidate, and this graphical analysis suggests that parties choosing a minority nominee in the primary may fare *better*, not worse, in the general election.

Table 1 presents estimates from several regression discontinuity specifications.<sup>4</sup> The first column shows a simple local linear approach that fits lines to the data on either side of the cut point within

<sup>4</sup>All tables include robust standard errors (Calonico, Cattaneo, and Farrell 2020; Calonico et al. 2021).

**Table 1.** RDD Estimates of the effect of nominating a minority candidate on general election vote share

	(1)	(2)	(3)	(4)
Minority nominee	0.07 (−0.01 to 0.16)	0.01 (−0.02 to 0.06)	0.03 (−0.01 to 0.08)	0.03 (−0.02 to 0.07)
Polynomial	p = 1	p = 1	p = 3	p = 1
Covariates		X	X	X
Effective sample size	215	129	264	97
Bandwidth	0.28	0.16	0.35	0.1
Bandwidth selection	Auto	Auto	Auto	Manual

an automatically selected bandwidth. The second adds district-level covariates,<sup>5</sup> which are not needed for identification but may increase the precision of the estimates. The third column uses a third-order local polynomial regression to model the data on either side of the cut point, again using an automatically selected bandwidth, and the fourth column uses a manual bandwidth of 0.1. These estimates vary slightly in size but yield similar overall conclusions: a slight positive effect and thus, if anything, an electoral *advantage* to parties that nominate the minority candidate in a cross-racial primary. While the positive effect is not significantly different from 0, we can rule out substantial electoral penalties: the 95 per cent confidence interval for the estimates in column 1, for example, excludes minority candidate electoral penalties of one percentage point or larger.

These estimates go against the conventional wisdom that parties will pay an electoral penalty should they nominate a minority candidate, at least in districts where there are closely fought cross-racial primaries. We examine the robustness of these findings in the Online Appendix. First, Section A3 explores many plausible RD specifications, which include bandwidth selection approaches, kernels, polynomial order, and bias-correction decisions. The vast majority of these specifications yield positive estimates, such as those shown in Fig. 1, and never document a significant negative effect. All of our evidence points against an electoral penalty when parties nominate minority candidates.

Second, we address whether the sample of elections used here is dominated by majority-minority or heavily Democratic districts where an electoral penalty is especially unlikely. Section A1 of the Online Appendix demonstrates that the districts in our sample, though more diverse and slightly more Democratic than average legislative districts, are not extreme cases. Fewer than one-quarter of the districts in the sample are ‘majority-minority’ districts, and most are competitive between the parties. An analysis limited to Republican primaries provides no stronger evidence of an electoral penalty.

Finally, we note that our estimates do not measure a causal effect of the candidate race itself but include the selection of a particular candidate in a given race (Grumbach and Sahn 2020; Hall 2015; Marshall 2021). The RD design does not change any individual candidate’s identity, it compares otherwise similar elections that yielded nominees of different races, a distinction we discuss in more detail in the Online Appendix (Grumbach and Sahn 2020). This means that candidate characteristics *other than race* could also vary across the cut point if these characteristics are *correlated* with race (Sen and Wasow 2016). We view this possibility not as a threat to inference but as an important part of interpreting these estimates. The value of using real-world data is that we can see how actual candidates fare in real elections, so if some characteristics tend to vary with race, we view that as an important way to draw out the real-world implications of our study. Accordingly, we briefly examine discontinuities in other candidate characteristics in Appendix Figure A6. We show that when parties narrowly nominate a candidate of colour, they choose a bundle of characteristics. Those narrowly winning minority candidates are more likely to be newcomers running for a seat. In the Online Appendix, we also discuss how these correlations reveal potential mechanisms by which the candidate’s race could affect election outcomes.

<sup>5</sup>Party, district population, an indicator for open-seat races, and the primary party’s vote share in the district in the 2016 presidential election.

Further, we consider threats to the interpretation of these estimates as a causal effect of nominating candidates of colour, such as the possibility of strategic discrimination by primary voters leading to ‘compensating differentials’ in candidate quality (Marshall 2021).

## Implications

Our unique data and design allow for a novel test of the ‘racial penalty’ story, which our results find no support for this perspective in recent state legislative elections. In this set of real-world elections, where parties narrowly nominate candidates of colour, parties do not face electoral penalties in the general election. Our evidence reinforces contemporary work that finds the same patterns using broader samples of legislative elections without restricting itself to close primaries (Fraga, Shah, and Juenke 2020; Juenke and Shah 2016). These results are also more in line with theoretical work on the partisan effects and implications of racial polarization, suggesting that an out-partisan racial penalty may be overwhelmed by a racial *advantage* from co-partisans (Agadjanian et al. 2020; Stout 2020; Tesler and Sears 2010). Though we are unable to unpack specific voter mechanisms here, our results should encourage more research into how voters use partisan shortcuts in combination with racial/ethnic cues to make both negative *and* positive evaluations of candidates for office (Chudy 2021; Crowder-Meyer, Gadarian, and Trounstein 2020a).

We hesitate to overgeneralize from these findings. However, in this brief note, we can situate the results within some of the larger conversations in the literature. Racial polarization in the parties, which has been documented since at least the late 1970s but has accelerated in the last ten to fifteen years (Enders and Scott 2019; Zhirkov and Valentino 2022), has co-occurred with asymmetrical descriptive representation in the parties. These two related patterns may help explain why partisanship can dominate voters’ electoral choices and, simultaneously, their racial *policy* preferences, resulting in increased descriptive representation with muted substantive policy success. The parties have become more clearly differentiated on racial politics; for example, criminal justice, immigration, voting access, and education policies, regardless of the identities of the party messengers.

Questions remain about the specific processes of candidate emergence that yield the elections studied here. We focus on a set of primaries where multiple candidates of various racial or ethnic backgrounds emerged, but not all districts currently see this sort of intra-party competition. Our results describe what happens in actual multiracial elections right now. Still, it would be worth revisiting these results if there were substantial shifts in party elites’ candidate recruitment practices or other election dynamics.

These results challenge the idea that parties consistently face an electoral penalty when they nominate racial and ethnic minority candidates (Doherty, Dowling, and Miller 2019; Fraga and Hassell 2021; Stephens-Dougan 2021). Such a perspective may have kept party elites from promoting and supporting candidates of colour, thus reinforcing the bigotry that undermines the fortunes of ambitious and talented candidates of colour. While racial bias is a fact of life in the American electorate, polarization has shifted the landscape of descriptive representation in the United States. We show that racial and ethnic minority candidates are viable and could potentially boost the electoral prospects of a party that chooses to nominate them.

**Supplementary Material.** The supplementary material for this article can be found at <https://doi.org/10.1017/S0007123423000583>

**Data availability statement.** Replication data for this article can be found in Harvard Dataverse at: <https://doi.org/10.7910/DVN/U5PDBR>.

**Acknowledgements.** Our thanks go to members of the Candidate Characteristics Cooperative (C3) for their work in gathering data on candidate demographics. We thank Kelisa Cain and Diana Camilla Valerio for their excellent research assistance. We also thank audience participants at APSA 2021 and the MSU American Politics workshop, as well as Loren Collingwood, Jennifer Wolak, Justin de Benedictis-Kessner, and Devin Caughey for insightful comments on early drafts of the paper. All errors are our responsibility.

**Financial support.** None.

**Competing interest.** None.

## References

- Agadjanian A et al.** (2020) Disfavor or Favor? Assessing the Valence of White Americans' Racial Attitudes. *Quarterly Journal of Political Science* **18**(1), 75–103
- Andersen DJ and Junn J** (2010) Deracializing Obama: White voters and the 2004 Illinois US Senate race. *American Politics Research* **38**(3), 443–70.
- Branton RP** (2009) The importance of race and ethnicity in congressional primary elections. *Political Research Quarterly* **62**(3), 459–73.
- Broockman DE** (2014) Do female politicians empower women to vote or run for office? A regression discontinuity approach. *Electoral Studies* **34**, 190–204.
- Bucchianeri P** (2018) Is running enough? Reconsidering the conventional wisdom about women candidates. *Political Behavior* **40**(2), 435–66.
- Calonico S, Cattaneo MD, and Farrell MH** (2020) Optimal bandwidth choice for robust bias-corrected inference in regression discontinuity designs. *The Econometrics Journal* **23**(2), 192–210.
- Calonico S et al.** (2021) *rdrobust: Robust Data-Driven Statistical Inference in Regression-Discontinuity Designs*. <https://cran.r-project.org/web/packages/rdrobust/rdrobust.pdf>.
- Canon DT and Posner RA** (1999) *Race, Redistricting, and Representation: The Unintended Consequences of Black Majority Districts*. Chicago, IL: University of Chicago Press.
- Chudy J** (2021) Racial sympathy and its political consequences. *The Journal of Politics* **83**(1), 122–36.
- Chudy J, Piston S, and Shipper J** (2019) Guilt by association: White collective guilt in American politics. *The Journal of Politics* **81**(3), 968–81.
- Citrin J, Green DP, and Sears DO** (1990) White reactions to black candidates: When does race matter? *Public Opinion Quarterly* **54**(1), 74–96.
- Crowder-Meyer M, Gadarian SK, and Trounstein J** (2020a) Voting can be hard, information helps. *Urban Affairs Review* **56**(1), 124–53.
- Crowder-Meyer M et al.** (2020b) A different kind of disadvantage: Candidate race, cognitive complexity, and voter choice. *Political Behavior* **42**(2), 509–30.
- De la Cuesta B and Imai K** (2016) Misunderstandings about the regression discontinuity design in the study of close elections. *Annual Review of Political Science* **19**, 375–96.
- Ditonto T** (2020) The mediating role of information search in the relationship between prejudice and voting behavior. *Political Psychology* **41**(1), 71–88.
- Doherty D, Dowling CM, and Miller MG** (2019) Do local party chairs think women and minority candidates can win? Evidence from a conjoint experiment. *Journal of Politics* **84**(4), 1282–97.
- Enders AM and Scott JS** (2019) The increasing racialization of American electoral politics, 1988–2016. *American Politics Research* **47**(2), 275–303.
- Fraga B and Hassell H** (2021) Are minority and women candidates penalized by party politics? Race, gender, and access to party support. *Political Research Quarterly* **74**(3), 540–55.
- Fraga BL, Shah P, and Juenke EG** (2020) Did women and candidates of color lead or ride the democratic wave in 2018? *PS: Political Science & Politics* **53**(3), 435–39.
- Fraga B, Juenke EG, and Shah P** (2021) Candidate Characteristics Cooperative (C3) 2018 Data.
- Grumbach Jacob M** (2022) Laboratories of democratic backsliding. *American Political Science Review* **117**(3), 967–84.
- Grumbach JM and Sahn A** (2020) Race and representation in campaign finance. *American Political Science Review* **114**(1), 206–21.
- Hall AB** (2015) What happens when extremists win primaries? *American Political Science Review* **109**(1), 18–42.
- Highton B** (2004) White voters and African American candidates for Congress. *Political Behavior* **26**(1), 1–25.
- Jackman S and Vavreck L** (2010) Obama's Advantage: Race, Partisanship, and Racial Attitudes in Context. Presented at the *Annual Meeting of the Midwest Political Science Ass'n*. Chicago, IL.
- Juenke EG** (2014) Ignorance is bias: The effect of Latino Losers on models of Latino representation. *American Journal of Political Science* **58**(3), 593–603.
- Juenke EG and Shah P** (2016) Demand and supply: Racial and ethnic minority candidates in White districts. *Journal of Race, Ethnicity, and Politics* **1**(1), 60–90.
- Kam CD** (2007) Implicit attitudes, explicit choices: When subliminal priming predicts candidate preference. *Political Behavior* **29**(3), 343–67.
- Kinder DR and Dale-Riddle A** (2012) *The End of Race?: Obama, 2008, and Racial Politics in America*. New Haven, CT: Yale University Press.

- Lewis-Beck MS, Tien C, and Nadeau R** (2010) Obama's missed landslide: A racial cost? *PS: Political Science & Politics* 43(1), 69–76.
- Lublin D** (1999) *The Paradox of Representation*. Princeton, NJ: Princeton University Press.
- Marshall J** (2021) Can close election regression discontinuity designs identify effects of winning politician characteristics? Working Paper.
- Sen M and Wasow O** (2016) Race as a 'bundle of sticks': Designs that estimate effects of seemingly immutable characteristics. *Annual Review of Political Science* 19, 499–522.
- Sigelman CK et al.** (1995) Black candidates, white voters: Understanding racial bias in political perceptions. *American Journal of Political Science*, 39(1), 243–65.
- Stephens-Dougan L** (2021) The persistence of racial cues and appeals in American elections. *Annual Review of Political Science* 24, 301–20.
- Stout CT** (2015) *Bringing Race Back In: Black Politicians, Deracialization, and Voting Behavior in the Age of Obama*. Charlottesville, VA: University of Virginia Press.
- Stout CT** (2020) *The Case for Identity Politics: Polarization, Demographic Change, and Racial Appeals*. Charlottesville, VA: University of Virginia Press.
- Tesler M** (2013) The return of old-fashioned racism to White Americans' partisan preferences in the early Obama era. *The Journal of Politics* 75(1), 110–23.
- Tesler M and Sears DO** (2010) *Obama's Race: The 2008 Election and the Dream of a Post Racial America*. Chicago, IL: University of Chicago Press.
- Valentino N, Hutchings V, and White I** (2002) Cues that matter: How political ads prime racial attitudes during campaigns. *American Political Science Review* 96(1), 75–90.
- Voss DS and Lublin D** (2001) Black incumbents, white districts: An appraisal of the 1996 congressional elections. *American Politics Research* 29(2), 141–82.
- White A et al.** (2023) Replication Data for: Evaluating the Minority Candidate Penalty with a Regression Discontinuity Approach. Available at <https://doi.org/10.7910/DVN/U5PDBR>, Harvard Dataverse, V1.
- Zhirkov K and Valentino NA** (2022) The origins and consequences of racialized schemas about U.S. Parties. *Journal of Race, Ethnicity, and Politics* 7(3), 484–504.