Local Queens and Princesses: The Impact of Female Dynasts on Provision of Public Goods

Aqila Putri*

Click here for the most recent version

Abstract

Women who reach political office often come from political dynasties, and the evidence on both female and dynastic politicians varies with the setting. I study what happens to public goods when Indonesian districts, which are responsible for basic services, elect a female dynast. I combine district elections from 2005 to 2020 with eight waves of a national village census, and I estimate the effect with a heterogeneity-robust event study estimator in a staggered difference-in-differences design. Electing a female dynast raises an index of public infrastructure by 0.075 of a village standard deviation, mostly through household infrastructure and education. The gain is concentrated among women from entrenched political families who did not take the seat directly from a relative, while men from the same families add almost nothing.

*University of Maryland, aputri@umd.edu. I am grateful for M.R. Sharan, Cory Smith, Ethan Kaplan, Erich Battistin, and Iman Sen for their guidance and encouragement throughout this project. I am grateful to Yoes Kenawas, whose dataset and earlier work on Indonesian political dynasties made this project possible. I also thank Ken Leonard, Josh Linn, Louis Preonas, Aichiro Suryo Prabowo, Michael Gibson, Juan Carlos Angulo, Fatima Najeeb, and Ariel Listo for the insightful comments on this paper. I gratefully acknowledge financial support from the Puri Family Dissertation Enhancement Award. All errors are my own.

1 Introduction

A politician's identity matters for the provision of public goods. The literature shows that female politicians implement more of the policies that women ask for. In West Bengal, India, villages where the council head seat was reserved for a woman received about nine more drinking water facilities and better roads, which are the public goods that women there raised with their councils more often than men did (Chattopadhyay and Duflo, 2004). Other dimensions of identity matter as well. When Indian states reserved legislative seats for scheduled castes and scheduled tribes, the policies aimed at those groups expanded. Reservation for scheduled castes raised job quotas for scheduled castes, and reservation for scheduled tribes raised welfare spending on scheduled tribes (Pande, 2003). At the same time, a growing body of literature shows that a larger share of female politicians than of male politicians comes from political dynasties, in developing and developed countries alike. In Indonesia, 44 percent of the women elected to the national parliament in 2019 came from political dynasties, against about a quarter of all members (Prihatini and Halimatusa'diyah, 2022). In Swedish municipal councils between 1998 and 2010, about 15 percent of elected women were dynastic, against about 11 percent of elected men (Folke, Rickne and Smith, 2021).

The prevalence of dynastic politicians among elected women raises the question of what they deliver once in office. The answer is not obvious because a female dynast brings both her family and her gender into office. Her family's network can add to her capacity to govern, but the family can also use the same network to capture the office for its own benefit (Querubin, 2016; Smith, 2012; Malik, Mirza and Platteau, 2021), and the empirical evidence on dynasties is mixed to negative (Bragança, Ferraz and Rios, 2015; George, 2020; Malik, Mirza and Platteau, 2021). As a woman, she may govern in her own right, or she may govern on her family's behalf (Labonne, Parsa and Querubin, 2021). The evidence on female politicians also differs across settings (Hessami and da Fonseca, 2020; Ferreira and Gyourko, 2014; Bagues and Campa, 2021).

Using Indonesian district-level data on elections and on public goods provision, I analyze what happens to the provision of public goods when citizens directly elect female dynasts as district heads. Indonesian districts are a good setting for this question because decentralization changed what they are responsible for. When Indonesia democratized in 1998 after three decades of authoritarian rule, it transferred a large share of the central government's authority to regional governments. Decentralization made district governments responsible for basic services, including education, health, public works, and housing. Importantly, district heads now propose the budget that pays for these public goods. This allows them to

set the policy agenda in their districts *and* creates many opportunities for local politicians to amass power (Government of Indonesia, 2014; Kis-Katos and Sjahrir, 2017).

I assemble a dataset that covers every district election in Indonesia from 2005, when direct local elections began, through 2020. For each election, the dataset records the winning head and deputy, their gender, their dynastic status, and their relatives who hold or have held public office. I define a dynast as a politician with a relative who holds or has held elected office. That relative's office can be anywhere in Indonesia and at any level of government, so a woman is a dynast under this definition whether or not her family has ever held the district she now leads. I count a district as treated when a female dynast wins either seat, because the deputy has a statutory role over the same service delivery. I combine the dynasty dataset with eight waves of an Indonesian village census from 2003 to 2021. Following Martinez-Bravo, Mukherjee and Stegmann (2017), I use the census to measure the provision of education, health, and household infrastructure. The estimation sample contains 1,741 elections across 514 districts, and 102 of these elections, in 76 districts, elected a female dynast.

I exploit the staggered timing of Indonesian district elections in a difference-in-differences design, and I estimate the impact of female dynasts on the provision of public goods with the heterogeneity-robust event study estimator of De Chaisemartin and d'Haultfoeuille (2020). I find that electing a female dynast raises the aggregate index of public infrastructure by 7.5 percent of a village standard deviation. Measured against the variation across districts and survey waves, the effect is about 13.5 percent of a standard deviation. The effect comes mostly from household infrastructure and education. Household infrastructure rises by 7.3 percent of a village standard deviation and education by 5.7 percent, while the effect on health is small. The gains in household infrastructure are concentrated in less urbanized districts, which had fewer household amenities to begin with.

To interpret these results, I build a simple model of a district head's payoff from office that incorporates both the family and the gender dimensions of her identity. In the model, a family's network can raise the public goods that a district head delivers from each unit of the budget, but it can also raise the private return to diverting the budget. Which of the two dominates is an empirical question. The model also predicts that the network helps a district head more when she places more weight on delivering public goods. If a woman who governs in her own right places more weight on delivery than a man, she should gain more from her family's network than a man does. A woman who governs on her family's behalf takes on the family's weight, so she should gain the same as a man.

The model implies that the effect of a dynasty depends on how much of a network the family brings into office, and dynastic families differ a lot in this respect. For instance,

families that first held public office before 2005, when district heads were chosen by the central government and then by the district parliament, have much broader networks than families that entered office later. In the elections that female dynasts won, the older families have on average 3.5 relatives in office, against 1.7 for newer families. They have 2.1 relatives in legislatures, against 0.5. They also have more relatives outside the district and at the provincial level. The family's network should therefore help some district heads much more than others.

To see how the data line up with the predictions of the model, I start with the average across all families. I evaluate whether the effect comes from her gender, from her dynastic status, or from both in a two-by-two matrix comparison. I estimate the effects of female dynasts, female non-dynasts, and male dynasts in one regression, with each group measured against a male non-dynastic district head. Holding dynastic status fixed, women raise the aggregate index by 3.7 percent of a standard deviation more than men. Holding gender fixed, dynasts raise it by 1.6 percent of a standard deviation more than non-dynasts. The interaction between gender and dynastic status is close to zero, so on average the effect of being a woman and the effect of being a dynast simply add up. Because this average pools families with broad networks and families with thin ones, an interaction close to zero is within what the model allows.

I further explore these dynamics in the mechanism analysis, following the structure of the model. First, the effect is concentrated in the families that first held public office before 2005, which I refer to as entrenched families from here on. Female dynasts from entrenched families raise the aggregate index by 9.4 percent of a standard deviation, against 2.6 percent for female dynasts from newer families. For female dynasts, then, the capacity that an entrenched family brings outweighs the capture it enables. Entrenched families are also no more likely than newer families to have held the seat that the female dynast now holds, and the effect is largest where her family never held the seat.

Second, entrenched families help women but do not help men. Within entrenched families, female dynasts raise the aggregate index by 5.8 percent of a standard deviation more than male dynasts. Male dynasts from entrenched families raise it by less than 1 percent of a standard deviation relative to a male non-dynastic district head. Third, the gain comes from women who are unlikely to be governing on their family's behalf. A woman who takes the seat directly from a relative is the most plausible case of a proxy, so I divide the female dynast tickets by whether they directly succeeded a relative. Among tickets that did not, female dynasts from entrenched families raise the aggregate index by 8.2 percent of a standard deviation more than male dynasts. In contrast, women who directly succeeded a relative are

indistinguishable from the men in their families.

Finally, the women who carry the effect are as experienced at their first election as men from entrenched families, but more experienced than other women. Most of this experience comes from careers in legislatures. Dynastic women also win elections as often as dynastic men. The gap between women and men from entrenched families therefore does not come from differences in formal experience or from a higher bar at the ballot box.

I also examine two alternative channels. In districts with higher female employment, which I use as a proxy for more favorable gender norms, the effect arrives sooner. Alignment with the provincial governor does not appear to be the channel. Most of the differences between groups in the mechanism analysis are not statistically significant. These results therefore establish an ordering of the estimates, and I interpret them as suggestive evidence.

This study contributes to four strands of literature: the impacts of female leaders, the mechanisms behind female politicians' governance, the impacts of dynastic politicians, and the intersection of gender and dynastic politics. A large literature has shown that female leaders provide more of the public goods that women value, engage less in rent-seeking behavior, and learn the ropes quickly on the job (Dollar, Fisman and Gatti, 2001; Chattopadhyay and Duflo, 2004; Bhalotra and Clots-Figueras, 2014; Brollo and Troiano, 2016; Afridi, Iversen and Sharan, 2017; Baskaran and Hessami, 2019; Lippmann, 2022; Bhalotra et al., 2023). These effects are concentrated in developing countries, and they are smaller or absent in developed countries (Hessami and da Fonseca, 2020; Ferreira and Gyourko, 2014). I contribute to this literature by showing that female politicians from dynastic families in a developing democracy raise the provision of public goods, particularly household infrastructure in less urbanized areas. How much they deliver depends on the family they come from.

This study also contributes to the literature on the mechanisms through which electing female politicians changes the quality of governance. Within this literature, Besley et al. (2017) show that a gender quota in Sweden pushed out mediocre male politicians and raised the competence of politicians overall. Andersen et al. (2008) show that, in a public goods game, men in a matrilineal society in India contribute more than men in a patrilineal society. These studies point to the competence of the politicians who are selected and to the norms of the society. I propose a mechanism that runs through the family. In my framework, a woman's weight on delivering public goods interacts with her family's network, and I show evidence consistent with it. An entrenched family's network helps women and does not help men. The gain comes from women who did not take the seat directly from a relative, and these women bring the same formal experience as men from the same families.

The literature on the economic impact of political dynasties has found mostly negative

or null effects, and these effects depend on the setting. Dynasties have been shown to slow economic growth, to spend more on public services without improving them, and to lower short-run economic performance, particularly where they are entrenched and electoral competition is close (Asako et al., 2015; Bragança, Ferraz and Rios, 2015; Dar, 2019; George, 2020; Malik, Mirza and Platteau, 2021). This study contributes to this literature by showing that the effect of a dynasty depends on the kind of family and on who governs. Entrenched families, whose networks are broad, help female dynasts deliver more public goods. The same families add almost nothing when the dynast is a man. This offers one reason why the effects of dynasties differ across settings.

Lastly, this paper contributes to the literature on the intersection of gender and dynastic politics. The most closely related work shows that female politicians are more likely than male politicians to be dynastic (Dal Bó, Dal Bó and Snyder, 2009; Geys and Smith, 2017; Smith and Martin, 2017), that parties and voters use a female politician's relative as a signal of her quality (Folke, Rickne and Smith, 2021), and that the female relatives of term-limited incumbents run for office to keep the seat in the family (Labonne, Parsa and Querubin, 2021). These studies focus on how female dynasts enter office. This study shows what they deliver once in office, and that what they deliver depends on whether they plausibly govern on their family's behalf. Women who take the seat directly from a relative deliver about what the men from their families deliver. Women from entrenched families who did not take the seat directly from a relative deliver more.

The remainder of the paper is organized as follows. Section 2 sets out the conceptual framework. Section 3 discusses the Indonesian context, and Section 4 describes the data. Section 5 sets out the empirical strategy. Section 6 reports the main results, and Section 6.3 separates gender from dynastic status. Section 7 reports heterogeneity, and Section 8 examines the mechanisms. Section 9 concludes.

2 Conceptual Framework

Why does a politician's identity matter for the public goods they deliver? When politicians cannot commit in advance to the platforms they campaign on, the person who is elected brings her own preferences and her own capacity into office with her, and the policy that results reflects both of those things (Osborne and Slivinski, 1996; Besley and Coate, 1997). This is what separates the citizen-candidate account from the median voter benchmark, in which only the preference of the median voter determines policy (Downs, 1957). It is also why the question of who wins is an economic question as well as a political one. By now

there is a large body of evidence showing that the identity of a politician moves policy, and a great deal of that evidence is about gender in particular (Chattopadhyay and Duflo, 2004; Hessami and da Fonseca, 2020).

What happens when a female dynast is elected is an empirical question. There are two dimensions that shape her preferences and her capacity, and within each one there are two forces that can pull in opposite directions. The first dimension is her family, and her family offers a potential contrast of capacity against capture. The second dimension is her identity as a woman, and it offers a contrast of own standing against proxy.

Having a family member who has held public office means that the family has ties into the governing apparatus: the bureaucracy, the parties, the legislative chamber, and the higher levels of government. She can then take advantage of this existing apparatus to lower the cost of governing: projects are approved, funded, and built more easily and in larger numbers. The family connection then adds to her *capacity* in governing.

However, the same family network can also serve as a *capture* mechanism for her family. Instead of using the apparatus as a boost to her policymaking, the family may use it to extract rents and to shield itself from electoral competition. This is consistent with the literature on the persistence of elite capture from the authoritarian era (Martinez-Bravo, Mukherjee and Stegmann, 2017), and with evidence that development outcomes are worse where entrenched dynasties win (Malik, Mirza and Platteau, 2021).

One common thread runs through both the capacity and the capture forces: both are strongest when the family has been entrenched in the system for a long time. Holding office makes it easier for a family to keep holding office (Querubin, 2016), and an old family has had time to build and foster its network and connections, which makes its advantage stronger than that of a newer family. Because both forces are strongest in old families, entrenchment is where the net effect of the family should be most visible. In Indonesia, the oldest families are those that held office before voters chose their district heads directly, a history I describe in Section 3.4.

The second dimension that affects a female dynast's preferences and capacity is her identity as a woman. Similar to the family network, this dimension can pull her governance in more than one direction. At the baseline, the literature has shown that women in public office provide more of the public goods that women value and extract fewer rents (Chattopadhyay and Duflo, 2004; Dollar, Fisman and Gatti, 2001; Brollo and Troiano, 2016). But a female dynast is tied to her family's network, and this creates a different calculus from that of a woman without a political family. The contrast created by this dimension is own standing against proxy, and it is a question of whether her gender shows up in how she governs. A

female dynast who governs in her own right should look like the baseline female politician. In this scenario, she builds her own legitimacy and uses the family's capacity to help her govern better and provide more public goods.

But she can instead serve as a proxy for her family. Instead of using the network to boost her own governing, she governs on behalf of a member of her family. In this case, the family's behavior carries on, whatever mix of capacity and capture it involves, and she becomes indistinguishable from a male member of her family. This is consistent with the "benchwarmer" logic of Labonne, Parsa and Querubin (2021), in which female relatives run for the seat of a term-limited incumbent. In Indonesia, a district head who has served two terms can never hold the same seat again, so the question is not who keeps the seat for him but who governs during her term.

Which forces win within each dimension, and how the two dimensions interact, can be illustrated with a simple model of a district head's payoff, which I set out in Appendix A. In the model, a district head controls a budget and decides how much of it to divert to rents for herself or her family. She spends the rest on public goods. Her family's network does two things at once. It raises the amount of public goods that each unit of the budget delivers, which is the capacity force, and it raises the private return to diverting the budget, which is the capture force. The district head also places a weight on delivering public goods. This weight captures her own preferences and the value she places on building a reputation of her own.

There are three results. First, the effect of the family network on public goods is ambiguous. A stronger network delivers more public goods for each unit of the budget, but it also makes diversion more attractive, and which of the two dominates is an empirical question. Second, the family network helps more when the district head places more weight on delivering public goods. A district head who cares more about delivery diverts less, so more of the capacity that her family brings reaches her constituents. Third, suppose a woman who governs in her own right places more weight on delivery than a man. Then she delivers more than a man from a family with the same network, and the gap grows as the network becomes stronger. A woman who governs on behalf of her family takes on the family's weight, and she delivers the same as a man from her family.

These results organize the empirical sections. The comparison between entrenched and newer families tells us which of the capacity and capture forces dominates for a female dynast. The comparison between female and male dynasts from entrenched families tells us whether the family's network helps a woman more than a man. The comparison between women who directly succeeded a relative and women who did not tells us whether the gain comes

from women who govern in their own right, because direct succession is the case in which governing on the family's behalf is most plausible.

3 Context

3.1 Decentralization and direct local elections

Indonesia was under an authoritarian regime for more than thirty years. The leader, Suharto, controlled the population and the opposition parties along with the military. In the wake of the Asian financial crisis that culminated in 1997, and after a series of corruption scandals that involved his family and his close circle, Suharto resigned from office in 1998 following an intense period of protests and riots. The 1998 event is commonly referred to as *reformasi*, and it marked the point at which Indonesia began the twin processes of democratization and decentralization, in which citizens played a far more central role in determining who would hold political office (Aspinall and Mas'udi, 2017). Elections with predetermined winners, which had been a feature of the Suharto regime, gave way to representative elections in which voters voted for a party, and the representatives of the parties that received enough votes to enter parliament then formed coalitions to select the president (Sherlock, 2004).

Importantly, decentralization moved authority and money downward. Decentralization devolved substantial spending and revenue authority to the autonomous tier of government below the province, *kabupaten* and *kotamadya*, which I refer to as *districts* throughout this paper.[1] Because the provinces received far less of it, the district became the level of government at which decisions about schools, clinics, and village infrastructure are actually taken and financed. When a village gets a new health post, or a road that a truck can still use in the rainy season, the decision and the money behind it come from the district. Almost overnight, the district headship became an office worth contesting, which makes the district the natural unit of analysis for this study. The same authority also makes the office worth capturing. In December 2016, for example, the Corruption Eradication Commission (KPK) arrested the *bupati* of Klaten in Central Java, Sri Hartini, for accepting bribes in exchange for positions in the district bureaucracy, and in 2017 the anti-corruption court sentenced her to eleven years in prison (CNN Indonesia, 2017). Her husband had been *bupati* of Klaten from

[1] Indonesia has two autonomous tiers of government below the national one. The province is the upper tier and the district is the lower one. A district takes one of two forms. A *kabupaten* is a district whose territory is largely rural and which contains villages, and a *kotamadya* is a city. The law gives the two the same responsibilities, and what separates them is the territory they cover and the administrative units beneath them.

2000 to 2005, and she had served as deputy from 2010 to 2015 before she won the headship herself (Tirto.id, 2016).

Under Suharto, district heads were in effect appointed by the central government through the Ministry of Home Affairs. After 1999 the local parliament elected them, and from 2005 citizens elected them directly. Direct elections for district head, which are known as *pilkada*, began in 2005. They were phased in as the last appointed heads were allowed to finish out their terms, which means that districts entered the electoral era at different dates and then ran on five-year cycles that were offset from one another. Beginning in 2015 the Election Commission started to synchronize election days within each cycle, with the goal that all districts will vote together by 2027. This history has two consequences for the research design. First, the timing of elections is staggered because of how decentralization was implemented—local incumbents did not choose it. Second, the earliest district-years in the data come before any direct election has been held in that district, and this is why the estimation sample begins when a district becomes at risk of treatment, which for most districts is later than the start of the outcome panel.

3.2 Responsibilities of a district head and deputy

A district is led by a *bupati* in a *kabupaten* and by a *walikota* in a *kotamadya*, and each of them serves alongside a deputy, the *wakil bupati* or the *wakil walikota*. Law 23 of 2014 on Regional Government (Government of Indonesia, 2014) assigns governmental affairs to *kabupaten* and *kotamadya* in the same terms throughout, so a *bupati* and a *walikota* answer for the same list of responsibilities. What differs between the two is the tier beneath them. A *kabupaten* contains *desa*, which are villages with their own elected heads and their own budgets, and a *kotamadya* contains *kelurahan*, which are administrative units of the city government.

This paper measures the affairs that the law makes mandatory for a district. Article 12 of Law 23 of 2014 lists six mandatory affairs tied to basic services: education, health, public works and spatial planning, public housing and settlements, public order and the protection of the community, and social affairs. Four of those six cover what the village census records, which is schools, health facilities and health staff, roads, and water and sanitation.

The head holds the powers that decide what a district builds. Under Article 65, the *bupati* or the *walikota* leads the implementation of the affairs that fall to the district, submits the long-term and the medium-term development plans to the local parliament, sets the annual work plan, and proposes the district budget together with its accountability report. The deputy's duties are set out in Article 66. The deputy assists the head in leading the

implementation of the district's affairs and in coordinating the district agencies. For a *wakil bupati* or a *wakil walikota* the law then assigns something more specific, which is monitoring and evaluating the administration carried out by the district agencies, the *kelurahan* and the *desa*. The deputy also exercises the head's duties and powers when the head is detained or is temporarily unable to serve, and the head may assign further duties to the deputy by decree.

3.3 Elite capture and the character of *pilkada*

Decentralization put resources into the hands of districts. At the same time, it opened up local political markets. Who led a district used to be decided from above, and now the district's residents choose who leads them. Many local elites used the arrival of electoral democracy as an opportunity to place family members and close associates into executive and legislative office (Aspinall and Mietzner, 2010), although Martinez-Bravo, Mukherjee and Stegmann (2017) find that districts which moved earlier to post-*reformasi* leadership have better governance outcomes than districts that were exposed for longer to Suharto-era appointees, and Hsiao (2019) finds that democratization reduced misallocation overall. Political scientists describe *pilkada* as an arena in which individual candidates are the key actors and political parties matter much less. The candidates who succeed tend to be locally powerful and wealthy individuals, and they run their campaigns through their own formal and informal networks (Aspinall and Mas'udi, 2017; Mietzner, 2020). Campaigns also rely a great deal on what is called money politics, in which candidates and parties give money, goods, and services to individual voters or to small groups instead of delivering common public goods such as health care or education (Aspinall, 2015). Wealth, networks, and clientelistic behavior are exactly the conditions under which family-based politics thrives.

3.4 Political dynasties in Indonesia

A dynastic politician is defined as an individual who shares a familial relationship—by blood or marriage—with at least one person who holds or has held an executive or legislative office (Geys and Smith, 2017). Dynasty formation in Indonesia predates *reformasi*, notably with the rise of the Sukarno clan. Sukarno, Indonesia's first president and founding father, was ousted by Suharto in the 1960s. The country later witnessed Megawati Sukarnoputri, one of Sukarno's daughters, rise as a beacon of resistance during Suharto's regime; she went on to become Indonesia's fifth president, from 2001 to 2004. The Sukarno dynasty has now reached its third generation, with Megawati's daughter Puan Maharani serving as Speaker of the People's Representative Council (DPR) (Mietzner, 2016). These dynasties exist well

beyond the national level. As post-*reformasi* electoral practice has shown, many local dynasties have emerged and thrived. One example is the dynasty of Ratu Atut Chosiyah, governor of Banten province from 2007 to 2014; during and after her time in office, her sister, stepbrother, stepmother, sister-in-law, and son held or currently hold executive positions across the province. National dynasties have likewise branched into the subnational level. Most prominently, the family of former president Joko Widodo has won executive offices across the country: his eldest son, Gibran Rakabuming Raka, won the mayoralty of Surakarta in Central Java—the same office Widodo himself held when he first entered politics—and in 2024 was elected Vice President of Indonesia alongside President Prabowo Subianto. His son-in-law won the mayoralty of Medan, North Sumatra, the largest metropolitan area outside Java.

The anecdotal evidence suggests that political dynasties in Indonesia have only grown over time. The Indonesian Ministry of Home Affairs estimated a 3% rate of dynastic growth in 2005, rising to 11% in 2013 (Prihatini and Halimatusa'diyah, 2022). The subsequent five years saw explosive growth: a 300% increase from 2013 to 2018, by which point 20% of provinces and districts were led by a dynastic politician (Kenawas, 2020). These figures capture only the executive branch; studies of dynasties in the legislature are not as robust as those for the executive. One recent study finds a 5% increase in female dynasts in the national legislature between 2009 and 2019 (Prihatini and Halimatusa'diyah, 2022). With multiple legislative bodies at the national, provincial, and sub-national levels, however, the true extent of dynastic politics may be substantially understated. Heading into the 2024 election, news coverage was saturated with anecdotal evidence of dynastic entrenchment—headlines reporting not only established dynasty names but also new family members of previously non-dynastic politicians entering the race. This proliferation of new dynasts makes Indonesia's political landscape a valuable setting for examining how the spread of political dynasties shapes people's livelihoods.

There is very little systematic information about the characteristics of the dynasts who hold executive office, even in the popular media, but there is some information available about the legislature. Prihatini and Halimatusa'diyah (2022) shows that while three of every five dynasts are men, around 44% of Indonesia's female parliamentarians are dynasts. Half of the parliamentarians who are thirty-five or younger are dynasts, while only about a fifth of middle-aged and older parliamentarians are. Once they are elected, the majority of the female dynasts in parliament are assigned to committees that handle issues coded as feminine and treated as low-prestige, such as health, manpower and transmigration affairs, education, youth, sports, and tourism.

Many of the families that hold district office today first entered public office before 2005,

when voters had no direct say over who held it. Under the New Order, the district parliament nominated three to five candidates for district head, sent at least two of them to the Minister of Home Affairs through the governor, and the Minister chose the head (Government of Indonesia, 1974). From 1999 to 2004, the district parliament elected the head and the deputy itself, and the President ratified the result (Government of Indonesia, 1999*a*). The members of the district parliament were in turn chosen by their parties. In the 1999 legislative election, voters chose a party, and the party decided which of its candidates took the seats it won (Government of Indonesia, 1999*b*). In 2004, a candidate could take a seat out of the party's order only by collecting enough votes on her own to fill a whole seat (Government of Indonesia, 2003). Only from 2009 did the number of votes a candidate received decide who sat in parliament (Mahkamah Konstitusi Republik Indonesia, 2008). Before 2005, then, the path into district office ran through the party leadership, the district parliament, and the central government. A family that held office in that period built its ties to these institutions before voters chose their district heads, and those ties did not disappear when direct elections arrived.

4 Data

4.1 Local elections and political dynasties

I assembled a dataset on Indonesia's direct local elections and on the dynastic status of the politicians who contest them. No single official source records both, so I built the dataset from Indonesian news reporting, from Constitutional Court decisions on election disputes, from local government records, from the official biographies of the candidates, and from the results that the General Elections Commission (KPU) published. I collected most of it by hand, and I then used an AI agent working under a written protocol to complete it. Appendix B describes how I assembled the data, cycle by cycle. The dataset covers the *pilkada* for every *kabupaten* and *kotamadya* from 2005, which is when direct local elections began, through 2024. For each election, it records the winning head and deputy, the gender of each of them, the dynastic status of each of them, the other candidate pairs who contested the election, and the share of the vote that each pair received.

There are a couple of difficulties in the data collection process. The first difficulty is that the early elections are poorly documented. The last district heads chosen before *reformasi* were allowed to finish their terms, so direct elections began in different districts at different times. The earliest rounds of *pilkada* were not recorded systematically at either the national

or the local level, which means that the results for those cycles do not exist in any single official source. For the 2005 to 2008 elections, I use the data from Martinez-Bravo, Mukherjee and Stegmann (2017) as the base, and for 2010 I use the results that the KPU published. I completed both by hand. There is no consolidated record for 2011, 2012 and 2013, so I compiled those elections myself through archival research in Indonesian online media, in Constitutional Court decisions, and in local government records. For the 2015, 2017, and 2018 rounds I build on the data that Kenawas (2022) collected and hand-coded. I am grateful to him for sharing it with me. For 2020 and 2024, I extracted the elections from the tables that the KPU published.

The second difficulty is that I cannot infer dynastic ties off of names. Indonesian naming does not generally follow the Western convention in which the first name is a given name and the last name is inherited. Parents and children frequently have different last names, and only some ethnic groups use clan names at all. This means that the surname-based method that Cruz, Labonne and Querubin (2017) use in the Philippines is not feasible in the Indonesian setting. I therefore establish every familial tie in the data from reporting, from official biographies, or from court and electoral records.

I define a dynastic politician as an individual who shares a familial relationship, either by blood or by marriage, with at least one person who holds or has held an executive or legislative office. This definition is the one commonly used in the literature (Geys and Smith, 2017). There are two things about it that matter later. The relative's office can be any elected executive or legislative office in Indonesia, at the district, the provincial or the national level. The relative's office also does not have to be the office that she herself wins. A woman is therefore a dynast under this definition even when nobody in her family has ever held the district she now leads. Following George (2020), I then separate dynastic founders from dynastic descendants. A founder is the first person in her family to hold office, and she becomes a dynast only once a member of her family is subsequently elected—a woman who is the first in her family to win a district headship is a founder on the day she wins it, and she becomes a dynast later, when her brother or her daughter wins something. A descendant is not the first person in her family to hold office, and she is already a dynast at the moment that she is elected. The treatment that I study in this paper is the election of a female dynastic *descendant*, and I choose the descendant definition because it means that every woman I compare already had the family tie at the time she was elected, while a founder's tie only comes into existence afterwards. Throughout the rest of the paper I use *female dynast* as shorthand for a *female dynastic descendant*. I specify when the distinction between founders and descendants matters.

Table 1 describes the district heads elected between 2005 and 2020 as individuals. It counts each elected head once per election, separately for women and for men, and it shows how many of each are dynasts, founders, and non-dynasts. Women are only about 6 percent of the district heads elected between 2005 and 2020, but close to two-thirds of them are dynasts, against about one in ten of the men.

Table 1: Elected district heads by gender and dynastic status, 2005–2020

	Heads	Share of heads (%)	Dynasts n (%)	Founders n (%)	Non-dynasts n (%)
Women	105	6.1	66 (62.9)	13 (12.4)	26 (24.8)
Men	1,629	93.9	164 (10.1)	352 (21.6)	1,113 (68.3)
All heads	1,734	100.0	230 (13.3)	365 (21.0)	1,139 (65.7)

There is one observation per elected district head per election, in *kabupaten* and *kotamadya*. A dynast is a dynastic descendant, who had a relative in elected office before her own election. A founder is the first person in her family to hold office, and she becomes a dynast when a relative is elected after her. Percentages in the last three columns are shares of the row. Heads whose gender is not coded are excluded.

Table 2 describes the elections in the data. I restrict the analysis to elections held in 2020 or earlier, because those are the elections for which I have at least one subsequent survey wave in which to measure an effect. This leaves me with 1,741 elections across 514 districts.

4.2 Relatives and kin in office

For each treated election, I take the relatives that the dynasty data records for the winning politician, and I match those relatives against a roster of the district's previous heads and deputies. This shows whether a member of the woman's family had held that district's executive office before her, and if one had, when. From this I construct two indicators. The first records whether her family had ever held the office before her. The second records whether her family held the office immediately before her.

There are two features of this matching procedure that matter for how the results in Section 8 should be interpreted. The first feature is that the classification fails much more easily in one direction than in the other. If I fail to match a relative to the roster, then I record a case where the family had no claim on the office when in fact it had one. A mistake in the other direction needs a name in the roster to be matched to a relative who is not the same person, and that is possible where names are common. The matching procedure guards against it, so I expect measurement error in this variable to run mostly in the first direction. I show in Section 8 what that means for the result I report there.

Table 2: Local elections, 2005–2020

Election year	Elections	Share of elections (%)			
		Woman on ticket	Dynast on ticket	Female dynast	Districts
2005	221	6.3	7.2	1.4	221
2006	76	11.8	7.9	2.6	76
2007	42	4.8	14.3	0.0	42
2008	143	2.8	14.0	0.7	143
2010	212	12.7	12.3	5.2	212
2011	67	4.5	9.0	0.0	67
2012	73	9.6	21.9	4.1	73
2013	132	12.1	20.5	4.5	132
2015	260	16.9	18.1	7.7	260
2017	100	9.0	18.0	4.0	100
2018	154	18.2	31.8	12.3	154
2020	261	21.1	33.0	12.6	261
All years	1741	12.5	18.6	5.9	514
Leader groups, all years (mutually exclusive and exhaustive)					
Female dynast	102		5.9%		76
Male dynast	221		12.7%		150
Female non-dynast	103		5.9%		80
Male non-dynast	1315		75.5%		493

There is one observation per district election. I count a ticket as female if the head, the deputy, or both are women, and as dynastic if the head, the deputy, or both are dynasts. In the leader groups, a female-dynast ticket has at least one woman who is a dynast. A male-dynast ticket has a dynast but no female dynast, a female non-dynast ticket has a woman but no dynast, and a male non-dynast ticket has neither.

The second feature is that the matching is imperfect. I encountered three main problems in the matching process. Indonesian officeholder records carry academic and honorific titles, they transliterate the same name inconsistently, and some of the names in them consist of a single token. All three of these defeat a naive matching procedure. A roster entry can carry a doctoral title in front of the name and two degree abbreviations after it while the dynasty record carries the bare name, and a procedure that compares the two strings directly will not report them as the same person. I resolved the cases that the procedure could not handle by hand, which means that part of the classification rests on judgment I exercised case by case. I record every one of those cases in the codebook that accompanies the data, which names the elections and makes them checkable.

I also build a measure of how broad each dynast's kinship network is. I follow the family

links outward, so that a relative of a relative also counts as kin, and I code each office-holding relative by the branch of government in which she serves and by the geographic tier at which she serves. This measure covers 666 dynasts, of whom 113 are women. I use it descriptively in Section 8.1.1, where I also describe its limitations.

Finally, for the winning heads and deputies who are female dynasts, male dynasts, or female non-dynasts, I collected biographical information: year of birth, education, the position each person held before the election, and the public offices each person held before. Appendix B.6 describes how I collected it.

4.3 Public goods: PODES

To measure the provision of public goods I use the Village Potential Statistics (*Potensi Desa*, which is commonly abbreviated as PODES). PODES is a census of Indonesian villages, and the Indonesian Central Bureau of Statistics conducts it. While PODES has the word "village" in its name, it does not refer only to small rural towns. It uses the word to denote the smallest unit of government, and that unit covers urban areas as well as rural ones. I use eight waves of PODES: 2003, 2006, 2008, 2011, 2014, 2018, 2020, and 2021. PODES records the stock of facilities that exist in a village. It does not record the services that people actually receive from those facilities. A district head's budget authority operates on the infrastructure stock, so the stock captures the immediate effect of her decisions, and I use it as the main outcome.

Table 3 reports the district-level mean of each index in each wave. The earlier waves contain fewer districts than the later ones, and this is because of district proliferation: the number of *kabupaten* and *kotamadya* in Indonesia grew over this period as existing districts split into new ones.

4.4 District budgets

I use the World Bank's INDO-DAPOER database for the fiscal evidence in Appendix G.2. This database reports the revenues and the expenditures of district governments annually from 2003 to 2020 for 508 districts. There are two features of this data that make it useful alongside PODES. First, the data is annual, where PODES is irregular and roughly triennial. This lets the mechanism analysis capture detail that the PODES data might miss. Second, it measures what a district government spends, where PODES measures what exists on the ground. A district head has direct authority over spending, so the budget data reflects that authority very closely.

Table 3: Index means by survey wave

Wave	Education	Health	Household	Aggregate	Districts
2003	0.069	-0.102	-0.654	-0.229	511
2006	0.049	-0.039	-0.600	-0.197	511
2008	0.163	0.092	-0.547	-0.100	511
2011	0.037	0.048	-0.317	-0.090	514
2014	0.140	0.165	-0.117	0.063	511
2018	0.156	0.231	0.250	0.212	514
2020	0.236	—	0.397	0.317	514
2021	0.214	0.255	0.435	0.301	514

Entries are district-level means of the standardized indices in each survey wave. I standardize the village-level indices using the mean and standard deviation of the 2011–2021 village sample, and then average them within a district and wave. The components of the health index are not collected in the 2020 wave. Earlier waves cover fewer districts because of district proliferation.

4.5 Construction of sample and treatment variable

I define a district as treated when it elects a female dynast as either the district head or the deputy. A district is treated in every survey wave during which a female dynast is serving as its head or deputy, and it is untreated in every wave during which none is. As I discuss in Section 3.2, the powers and responsibilities of the head and the deputy are closely tied to each other, so the decisions that a district head takes may be affected by her deputy. There are limits to what this justifies. The deputy is accountable to the head, and the deputy holds no independent power to propose a regulation or a budget. Thus, the case for counting a district as treated when a female dynast wins either seat is that the deputy carries a statutory role over the same service delivery, and it is not that the two offices hold the same authority. I estimate the head and the deputy as separate arms in Section 7 to check whether their effects differ.

Table 4 sets out the restrictions that take me from the full election data to the estimation sample. I drop the elections held after 2020, because the last PODES wave is 2021 and later elections have no survey wave after them. There are 102 treated elections across 76 districts. Of these elections, 66 elect a female dynast as head, 37 elect one as deputy, and one elects a female dynast to both positions at once. Table 5 reports the treated elections by year and by the position the woman won. A district can elect a female dynast in one cycle and then fail to elect one in the next, which means that treatment in this setting is not absorbing. A district that elects a female dynast in two separate cycles contributes two treated elections.

Table 4: Sample construction

	Elections	Districts
All district elections in the data	2249	514
Held in 2020 or earlier	1741	514
Electing a female dynast	**102**	**76**

Table 5: Treated elections by year and position

Election year	Head	Deputy	Both	Elections
2005	1	2	0	3
2006	0	2	0	2
2008	1	0	0	1
2010	7	4	0	11
2012	3	0	0	3
2013	4	2	0	6
2015	15	6	1	20
2017	3	1	0	4
2018	11	8	0	19
2020	21	12	0	33
Total	66	37	1	102
Distinct districts				76

A district appears once for each treated election, so a district that elects a female dynast in two separate cycles contributes two rows.

4.6 Construction of the outcome variables

I construct the outcome variables by following the approach of Martinez-Bravo, Mukherjee and Stegmann (2017). The construction runs in four steps, and the order of the steps determines how the coefficients should be read. First, I build three counts for every village in every wave: education, health, and household counts. The education count is the number of school buildings in the village at each level of schooling: kindergarten, primary, middle, and upper secondary. The count includes public and private schools, and Islamic schools, at every level. The health count is the number of doctors and midwives in the village. The household count is the number of the five household amenities the village has. The components are binary variables that record whether the majority of households in a village have access to safe drinking water, whether they have proper garbage disposal, whether they have a toilet either in the house or in the village, and whether they use a modern fuel for cooking. The

index also includes a variable that records whether the village has a road that four-wheeled vehicles can use throughout the year. I sum the five variables into a count that runs from zero to five.

Second, I turn each of the three counts into a z-score. I compute one mean and one standard deviation per count from the pooled village sample of the 2011 to 2021 waves, and I apply that same pair of moments to all eight waves. The moments come from the village distribution. The standard deviation of an index across districts is much smaller than its standard deviation across villages, so scaling by the district standard deviation would inflate every coefficient in the paper. A coefficient of 0.05 in the tables below therefore means five percent of a village standard deviation.

Third, I aggregate the standardized indices to the district level. The treatment is defined at the district level, so I take the mean across all the villages in a district in a given survey wave. A z-score is linear in the count underneath it, so the district mean of the village z-scores is the district mean of the count expressed in village standard deviations.

Fourth, I take the simple mean of the three district-level indices to form the aggregate index. Every index enters the mean with equal weight.

There are three components of the original indices in Martinez-Bravo, Mukherjee and Stegmann (2017) that I do not include, because PODES does not collect them consistently across the waves that I use. The first is the variable that records whether births are attended by a traditional birth attendant. The second and the third are the numbers of *puskesmas* and *polindes*. A *puskesmas* is a publicly staffed clinic that provides primary care, and a *polindes* is a village maternity post. A *puskesmas* can sit at either the sub-district or the village level, and larger and more urban areas usually have one at both levels. The 2008 wave records only whether a village has a *puskesmas* and not how many it has, and *polindes* are not recorded in every wave. The health index therefore counts only health staff, which are doctors and midwives. Several of the variables that make up the health index are not collected in the 2020 wave, so I measure health in seven of the eight waves.

5 Methodology and Empirical Strategy

The election of female dynasts is staggered over time and is not random, so I use an event study design to estimate the impact of female dynasts. I also report the static two-way fixed effects estimate alongside the main results for comparison.

5.1 Estimator and specification

Let g index districts and let t index the eight PODES waves, so that $t = 1$ is the 2003 wave and $t = 8$ is the 2021 wave. Let $Y_{g,t}$ denote one of the four indices in district g and wave t, and let N_g denote the number of villages in district g. Treatment is

$$D_{g,t} = \begin{cases} 1 & \text{if a female dynast is serving as head/deputy of district } g \text{ in wave } t, \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

One way to estimate an effect in a panel of this shape is the static two-way fixed effects regression,

$$Y_{g,t} = \alpha_g + \lambda_t + \beta D_{g,t} + \varepsilon_{g,t}, \tag{2}$$

where α_g is a district fixed effect, λ_t is a survey wave fixed effect, and β is the coefficient of interest. I estimate (2) weighting each district by N_g and clustering on the district. This is the estimate that I report alongside the robust one in every main table.

Districts in Indonesia hold their elections on staggered cycles because pre-*reformasi* district leaders, who were assigned to office instead of elected, were allowed to finish out their term (Martinez-Bravo, Mukherjee and Stegmann, 2017). This means that a female dynast takes office at a different time in each district. When the timing of treatment varies in this way, β in (2) does not recover a simple average of treatment effects. It recovers instead a variance-weighted average of all the two-by-two comparisons that the panel makes available, and some of those comparisons use a district that is already treated as a control for a district that is treated later.

To make that concrete, a district that elected a female dynast in 2010 would serve as a control for a district that elects one in 2018, even though the first district has had such a leader in office for eight years by the time the comparison is made. If treatment effects differ across cohorts, or if they change over time within a cohort, then some of these comparisons enter the average with a negative weight (De Chaisemartin and d'Haultfoeuille, 2020; Borusyak, Jaravel and Spiess, 2023; Callaway and Sant'Anna, 2021; Sun and Abraham, 2021). Jakiela (2021) provides a simple way to diagnose this problem directly. For these reasons I use a heterogeneity-robust estimator for every main estimate in this paper. I report the static two-way fixed effects estimate next to it in the main table for comparison, and I rely on the static estimator on its own only in Section 6.3, where the robust estimator cannot implement the design that I need, and as a check in Section 6.5.

5.1.1 Event studies estimator

I estimate the effect of electing a female dynast using the estimator of De Chaisemartin and d'Haultfoeuille (2020), in its dynamic version.[2] The estimator works in the following way. It compares districts that change treatment status between two consecutive periods with districts that have not yet changed status, it takes that comparison in first differences, and it then aggregates these switcher-level comparisons into an effect at each horizon after the switch. Because the control group at every horizon consists only of districts that have not yet switched, the estimator never uses an already-treated district as a control, and this is what removes the negative weights that I describe above. In the example above, the district that switches in 2018 is compared with the districts that had still not elected a female dynast by 2018, and with nobody else.

In detail, the estimator does the following. It indexes event time to the first wave in which a district's treatment differs from its wave-one treatment,

$$F_g \;=\; \min\{\, t : D_{g,t} \neq D_{g,1} \,\}. \tag{3}$$

The first wave in the panel is 2003, which comes before the first direct election for district head, so $D_{g,1} = 0$ for every district in the sample, and (3) reduces to the first wave in which the district has a female dynast in office. A district that never elects one has $F_g = \infty$ and enters the estimation only as a control.

Fix a horizon ℓ. The control group for switcher g at that horizon is the set of districts whose treatment has still not changed by the time g has been treated for ℓ waves,

$$\mathcal{C}_\ell(g) \;=\; \{\, g' : F_{g'} > F_g - 1 + \ell \,\}. \tag{4}$$

The building block is a difference in differences between g and that control set, taken from the wave before g switches to the wave ℓ periods after it switches,

$$\mathrm{DID}_{g,\ell} \;=\; \big(Y_{g,\,F_g-1+\ell} - Y_{g,\,F_g-1}\big) \;-\; \sum_{g' \in \mathcal{C}_\ell(g)} \omega_{g'} \big(Y_{g',\,F_g-1+\ell} - Y_{g',\,F_g-1}\big), \qquad \omega_{g'} = \frac{N_{g'}}{\sum_{h \in \mathcal{C}_\ell(g)} N_h}. \tag{5}$$

Let $\mathcal{S}_\ell$ be the set of switchers still observed at horizon ℓ. The horizon-ℓ effect is the average of (5) across them,

$$\widehat{\delta}_\ell \;=\; \sum_{g \in \mathcal{S}_\ell} \frac{N_g}{\sum_{h \in \mathcal{S}_\ell} N_h} \mathrm{DID}_{g,\ell}, \qquad \ell = 1, \ldots, 4, \tag{6}$$

and the cardinality of $\mathcal{S}_\ell$ is the number of switching districts that I report at the first and the fourth horizon in every table. The placebo at horizon ℓ is the same object run backwards

[2] I estimate it with the `did_multiplegt_dyn` command in Stata.

from the same base period, so it compares g with its controls over ℓ waves before g is ever treated,

$$\widehat{\delta}^{\text{pl}}_{\ell} \;=\; \sum_{g \in \mathcal{S}^{\text{pl}}_{\ell}} \frac{N_g}{\sum_h N_h} \left[\left(Y_{g,\,F_g-1-\ell} - Y_{g,\,F_g-1} \right) - \sum_{g' \in \mathcal{C}_\ell(g)} \omega_{g'} \left(Y_{g',\,F_g-1-\ell} - Y_{g',\,F_g-1} \right) \right], \qquad \ell = 1, 2, 3. \tag{7}$$

Under parallel trends and no anticipation, each $\widehat{\delta}^{\text{pl}}_{\ell}$ has expectation zero, and that is what makes the placebos a test of the design. The number that I report in every table is an average of the horizon effects,

$$\widehat{\delta} \;=\; \sum_{\ell=1}^{4} w_\ell \, \widehat{\delta}_\ell, \tag{8}$$

with weights w_ℓ that fall as the horizon lengthens, because fewer districts are still observed four waves after they switch than are observed one wave after. The early horizons therefore carry most of the reported average.

Equation (4) is what solves the problem in (2). The control group at horizon ℓ contains only districts whose treatment has not yet changed, so a district that has already elected a female dynast is never used as a control for a district that elects one later. That is the comparison which generates the negative weights, and (4) removes it by construction.

The estimation panel is a district-by-wave panel that uses all eight PODES waves. I index time by survey round, because the waves are not evenly spaced in calendar time: they fall in 2003, 2006, 2008, 2011, 2014, 2018, 2020, and 2021. This affects how the horizons should be read. One round after the election is three years for a district that is treated before the gap between 2003 and 2006, and it is one year for a district that is treated before the gap between 2020 and 2021. The horizons are therefore ordered correctly in event time, but they do not correspond to the same calendar distance for every cohort.

I estimate four effect horizons and three placebo horizons. I weight each district by the number of villages it contains, so that every village receives the same weight in the estimation. I cluster the standard errors at the district level because treatment is assigned at the district level. I include a district in the estimation sample once direct elections for district head have begun in that district.

5.1.2 Reading the results

The estimate at horizon ℓ from the robust event study estimator is a cumulative effect, and it accumulates ℓ periods of exposure to treatment. A path that rises across horizons therefore does not by itself show an effect that grows stronger the longer a female dynast stays in office.

I therefore plot the event study for the main results twice: the first shows the raw path, and the second normalizes the horizon coefficients by the doses that districts actually received, which gives the effect of one additional period of exposure.

I label the leads and the lags by event time. Event time t is the last period before a district first elects a female dynast, which is a survey wave in the main panel and a budget year in the fiscal panel. The lead at $t-3$ is therefore the placebo coefficient three periods before that reference period, and the lag at $t+1$ is the effect in the first period after the election.

When a table reports a single estimate, that estimate is the robust estimator's average cumulative effect per treated district unless stated otherwise. It is a switcher-weighted average of the cumulative effects at each horizon, which is one of the quantities that the estimator reports directly. Note that it is not a sum across horizons, nor is it a per-period rate. The average is identical under both versions of the event study, so the single estimate I report does not depend on which of the two per-horizon estimates I use.

5.2 Balance

Table 6 reports the comparison between treated and untreated districts using the 2003 PODES wave. The balance table uses the 2003 wave because it comes before the first direct election for district head, which means that it is a pre-treatment wave for every district in the sample. Any later wave would be pre-treatment for some districts and post-treatment for others since some districts had their first local election as early as 2005. I include island fixed effects in the comparison as Indonesia is made up of more than seventeen thousand islands with very different population densities. I account for this because population density is related to public-goods provision in both directions: denser places need more infrastructure, and more infrastructure attracts people to move in. Each district contributes a single observation in the 2003 wave so I use heteroskedasticity-robust standard errors.

6 Results and Discussion

6.1 Main results

Figure 1 shows the results from the robust event study estimator. The estimator produces an effect at each horizon after a district first elects a female dynast. In each panel, I show three leads and four lags. Each point is the cumulative effect of the treatment path since the district first elected a female dynast, so the effect at horizon ℓ accumulates ℓ periods of

Table 6: Balance in the 2003 survey wave

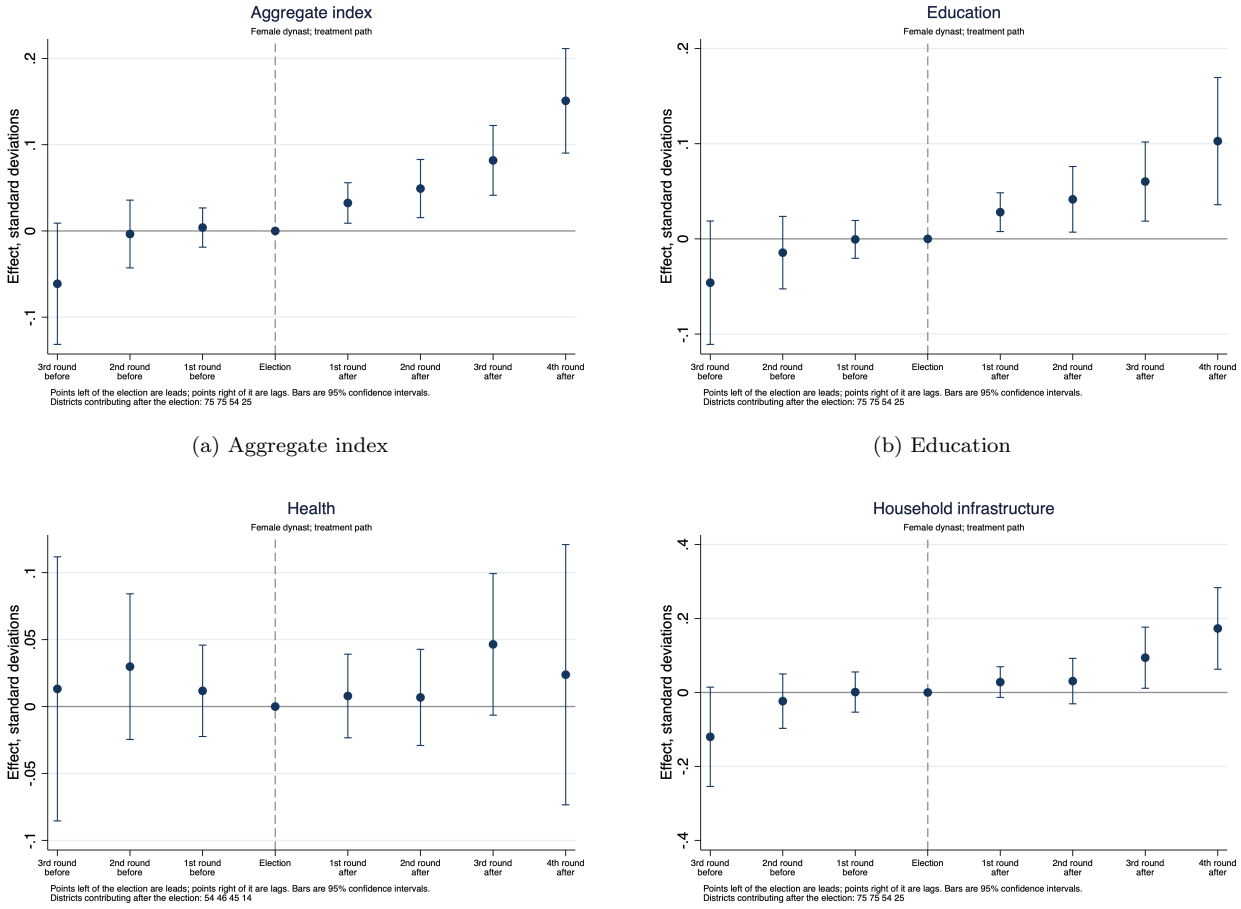
	Never treated		Ever treated		Within-island
	Mean	N	Mean	N	difference (SE)
Education index	0.047	362	0.189	67	-0.137
					(0.102)
Health index	-0.105	362	-0.087	67	-0.083
					(0.083)
Household index	-0.665	362	-0.593	67	-0.115
					(0.116)
Aggregate index	-0.241	362	-0.164	67	-0.112
					(0.092)
Villages in district	207.914	362	196.224	67	-8.522
					(15.863)
Joint F-test on the component indices: F = 1.50, p = 0.200					

Entries are wave-2003 means for each group and the difference between the groups within island, estimated by regressing each variable on a treatment indicator and island fixed effects. The within-island difference can have the opposite sign to the difference in means, because treated districts are concentrated on islands with higher values of the indices. Treated districts are those that elect a female dynast at any point in the sample window. The 2003 wave comes before the first direct election for district head, so it is pre-treatment for every district. Island fixed effects throughout. Each district contributes one observation, so the standard errors on the difference are heteroskedasticity-robust. Significance: *** 1%, ** 5%, * 10%.

exposure.

The aggregate and education panels rise steadily and monotonically. The aggregate runs from 0.032 at the first horizon to 0.151 at the fourth, and education runs from 0.028 to 0.103. In both cases every post-election coefficient is positive and every one has a confidence interval that excludes zero. Household infrastructure behaves differently. Its first two coefficients are 0.028 and 0.031, and their confidence intervals contain zero. The third and the fourth, at 0.094 and 0.173, do exclude zero, and they rest on 54 and 25 switching districts. The per-horizon estimates for health do not create any particular trend. Its coefficients are 0.008, 0.007, 0.046 and 0.024, with no ordering across horizons and confidence intervals that contain zero throughout. All four outcomes rest on 75 switching districts at the first horizon and on 25 at the fourth, except health, which runs from 54 down to 14. The largest coefficient in every panel is also the one with the thinnest support, and the confidence intervals widen accordingly.

Table 7 reports the effect of electing a female dynast on each of the four outcomes. I show the static two-way fixed effects estimate and the average cumulative effect per treated district of the robust event study estimator side by side, together with the two joint tests



(a) Aggregate index

(b) Education

(c) Health

(d) Household infrastructure

Figure 1: Event studies, cumulative effect by horizon

Each point is the cumulative effect of the treatment path since the district first elected a female dynast, so the effect at horizon ℓ accumulates ℓ periods of exposure. Coefficients at negative horizons are leads. Bars are 95 percent confidence intervals, clustered on district.

and the number of switching districts at the first and the fourth horizon. The standard error is clustered at the district level for both estimators.

Electing a female dynast raises the aggregate index by 0.0747 of a standard deviation, with a standard error of 0.0175. To give a sense of the size, the district-level aggregate index rose by about 0.087 per survey wave across the country between 2003 and 2021, so the effect is roughly one wave of average national progress. Two of the three components carry the aggregate. Household infrastructure is the larger of them at 0.0728 with a standard error of 0.0334, and education follows at 0.0574 with a standard error of 0.0190. Education has the best-behaved placebos of the four outcomes at a joint placebo p-value of 0.526. The village-level standard deviation of the school count is about five buildings, so an effect of 0.0574 of a standard deviation is on the order of a third of a school building per village. In

Table 7: Effect of electing a female dynast on public-goods provision

Outcome	Static TWFE	Robust estimator	Joint tests effects / placebo	Switchers h_1 / h_4
Aggregate index	0.0460***	0.0747***	0.000 / 0.227	75 / 25
	(0.0177)	(0.0175)		
Education	0.0403*	0.0574***	0.035 / 0.526	75 / 25
	(0.0217)	(0.0190)		
Health	0.0132	0.0222	0.356 / 0.678	54 / 14
	(0.0216)	(0.0191)		
Household infrastructure	0.0583*	0.0728**	0.008 / 0.279	75 / 25
	(0.0320)	(0.0334)		

The static estimate is β from equation (2). The robust estimate is the average cumulative effect per treated district, which is a switcher-weighted average of the horizon-specific cumulative effects and is not a sum across horizons. Estimates come from the heterogeneity-robust estimator of De Chaisemartin and d'Haultfoeuille (2020), on the district-by-wave panel, with districts weighted by their village count. Standard errors, in parentheses beneath each estimate, are clustered on district. Significance: *** 1%, ** 5%, * 10%. The two joint tests are nullity tests on the full set of horizon coefficients, so neither is a test of the estimate above it, and the two can disagree. The last column gives the number of switching districts at the first and the fourth horizon.

a district of two hundred villages, that is roughly sixty additional school buildings.

The estimate for health is the smallest of the four and it is not statistically significant. The estimate is 0.0222 with a standard error of 0.0191, and the joint test on the effects does not reject that the horizon coefficients are all zero. Health also has the thinnest support of the four outcomes at 54 switching districts at the first horizon and 14 at the fourth. This is largely because its components were not collected in the 2020 wave. Household infrastructure is significant at the five percent level, at 0.0728 with a standard error of 0.0334, and its joint test on the effects rejects at a p-value of 0.008. It is also the outcome that the heterogeneity and mechanism sections have the most to say about. The static estimates are smaller than the robust ones for all four outcomes. Among the three component outcomes the two estimators put them in the same order, with health the smallest and household infrastructure the largest. Table 8 restates the four estimates in raw units and as a share of the standard deviation of each index across districts.

6.2 The effect of one more period in office

The estimates above are cumulative, so a rising path shows exposure accumulating and does not by itself show an effect that strengthens the longer a female dynast stays in office. To separate the two, I normalize each horizon coefficient by the treatment doses that districts actually received. Because treatment here is binary, one dose is one survey round with a

Table 8: Magnitude of the estimated effects

Outcome	In raw units			Across districts	
	Unit	SD	Effect	SD	Effect (% of SD)
Aggregate index	—	—	—	0.56	13.5
Education	schools per village	5.02	0.29	0.65	8.8
Health	staff per village	4.91	0.11	0.42	5.3
Household	items per village	1.19	0.09	0.80	9.1

These are the estimates from Table 7 expressed on two other scales. The raw-unit columns give the village-level standard deviation of the underlying count and the estimate converted into that unit; the aggregate index is the mean of three z-scores measured in different units, so it has no raw unit. The across-district columns give the village-weighted standard deviation of each index across districts in the estimation sample, and the estimate as a percentage of it.

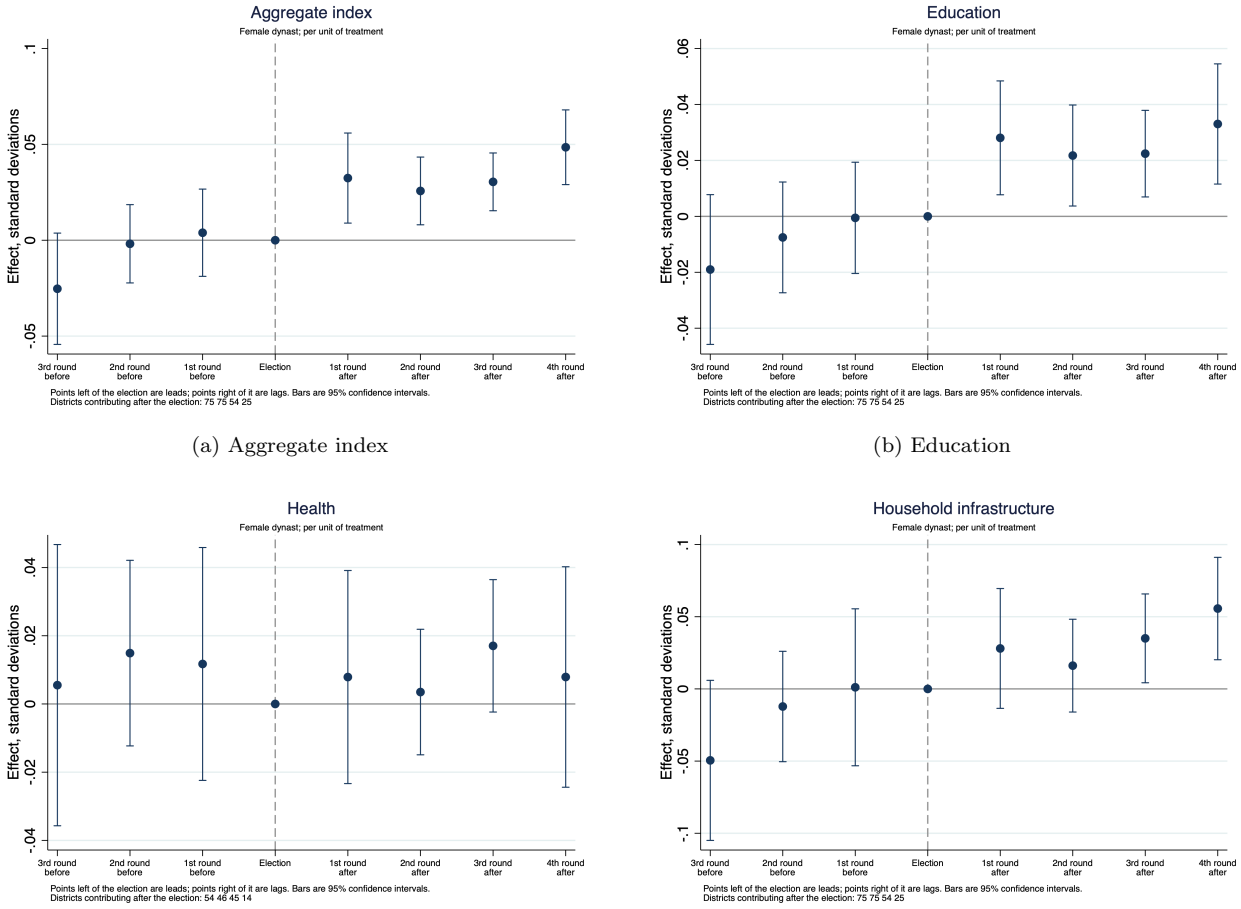
female dynast in office. This means that the normalized coefficient is equal to the effect of one additional such round. Figure 2 plots the four outcomes on that scale.

Once the estimates are normalized, the aggregate and education paths flatten almost completely, so almost all of the apparent intensification in Figure 1 is accumulating exposure. A district gets roughly the same improvement out of each additional round with a female dynast in office as it got out of the first one, and the growth in Figure 1 comes from the rounds adding up.

Household infrastructure is the exception. It is the outcome whose normalized path rises the most. It runs 0.028, 0.016, 0.035 and 0.056, so the fourth round delivers about twice what the first did. However, the support rests on 25 switching districts at that horizon and its confidence interval is wide. The lead at $t-3$ is -0.0495 with a t of -1.75, so I do not put much weight on the increasing path. Finally, the estimates for health are flat and small on this scale as well, at 0.008, 0.004, 0.017 and 0.008.

6.3 Specificity test

The results above are about a woman who is also a dynast, but they do not yet tell us which of the two characteristics is . To separate them, I sort every election into one of four mutually exclusive and exhaustive groups: a female dynast on the ticket (female dynast), a dynast but no female dynast (male dynast), a woman but no dynast (female non-dynast), and neither (male non-dynast). The four groups cover all 1,741 elections. I then look at those groups in two ways—individual regressions where each of the groups acts as the treatment, and one regression that includes three of the groups, omitting the fourth because of collinearity.



(a) Aggregate index

(b) Education

(c) Health

(d) Household infrastructure

Figure 2: Event studies, effect per period of exposure

Each horizon coefficient is divided by the treatment doses the switching districts actually received, so a point is the effect of one additional survey round with a female dynast in office. The first horizon matches the path figure, because one round of treatment is one dose. Bars are 95 percent confidence intervals, clustered on district.

6.3.1 What happens when each kind of leader is elected

The first way asks what happens when a district elects each kind of leader, taking one kind at a time as the treatment and every other district as the control. This gives me a level for each kind of leader on the same scale as the main result, and I estimate each one with the same robust estimator on the same panel. Figure 3 plots them and Table 9 gives the numbers.

Female dynasts are the only category whose joint effects test rejects, at $p < 0.001$. Male dynasts produce an estimate of 0.0286, which is about forty percent of the female dynastic effect, and its joint effects test does not reject, at a p-value of 0.256. Women who are not dynasts produce -0.0083, a point estimate essentially at zero. Pooling all dynasts regardless of gender produces an estimate of 0.0523, and that number is a mixture of both female

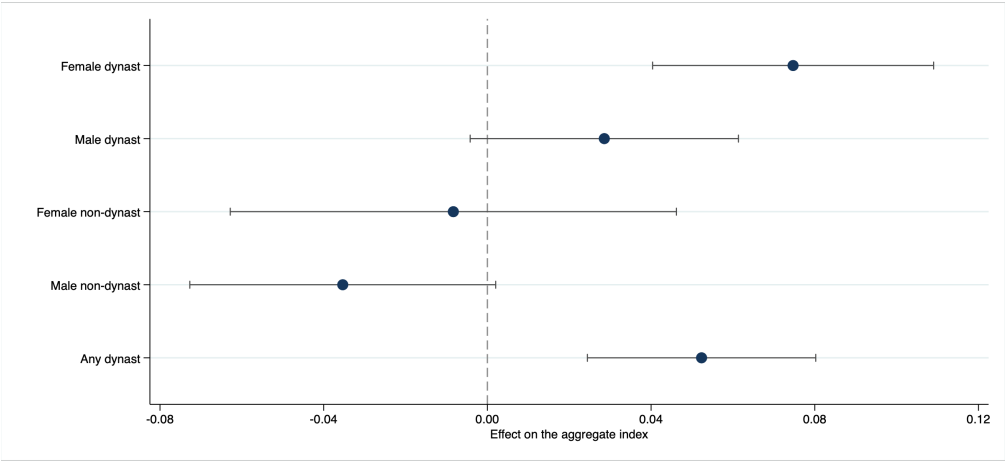


Figure 3: Effects by kind of leader elected

Each point is a separate estimation in which that kind of leader is the treatment and every district not electing that kind of leader is the control. Estimates come from the heterogeneity-robust estimator of De Chaisemartin and d'Haultfoeuille (2020), on the district-by-wave panel, with districts weighted by their village count. Bars are 95 percent confidence intervals, clustered on district. The bottom category pools the first two, so it is not mutually exclusive with them.

Table 9: Effects by kind of leader elected

Leader elected	Estimate	Joint p effects / placebo	Switchers h_1 / h_4
Female dynast	0.0747***	0.000 / 0.227	75 / 25
	(0.0175)		
Male dynast	0.0286*	0.256 / 0.964	118 / 49
	(0.0167)		
Female non-dynast	-0.0083	0.827 / 0.122	69 / 33
	(0.0278)		
Male non-dynast	-0.0353*	0.069 / 0.398	239 / 232
	(0.0191)		
Any dynast	0.0523***	0.001 / 0.417	177 / 71
	(0.0142)		

Each row is a separate estimation on the aggregate index, with that kind of leader as the treatment and all districts not in that category as the control. Entries are the average cumulative effect per treated district. Estimates come from the heterogeneity-robust estimator of De Chaisemartin and d'Haultfoeuille (2020), on the district-by-wave panel, with districts weighted by their village count. Standard errors, in parentheses beneath each estimate, are clustered on district. Significance: *** 1%, ** 5%, * 10%. The joint tests are nullity tests on the horizon coefficients, so neither is a test of the estimate. The final row pools the first two and is not mutually exclusive with them. This table reports levels; whether two categories differ from each other is not tested here.

and male dynasts' estimates. Finally, male non-dynastic leaders come in at -0.0353 with a p-value of 0.069 on the joint effects test. However, I do not interpret this as evidence that electing an ordinary man makes a district worse. That category contains 239 switching

districts against 75 for female dynasts, so it is close to the whole sample, and its control group is made up of the other three categories. A negative coefficient there is largely the arithmetic mirror of the positive coefficients above it.

Note that districts that elect a female dynast sit inside the control pool of every other row, which biases those rows downward. I measured that bias by dropping female-dynast districts from the control pool and estimating the rows again. Male dynasts move from 0.0286 to 0.0374 and male non-dynasts move from −0.0353 to −0.0455. The movement is in the predicted direction, but it does not change the interpretation of Figure 3.

6.3.2 Gender against dynastic status

The second way puts all four groups into one regression, which gives me the contrasts between gender and dynasty network directly. I estimate this with the static two-way fixed effects estimator instead of the robust event study estimator. This is because the robust estimator only takes a single treatment variable and cannot implement a four-group design. Note that the four indicators sum to one in every election, and the only survey wave in which all four are zero is 2003. This means that entering all four would make them a linear combination of the year dummies. I therefore omit the group with no woman and no dynastic descendant, which means every coefficient below is measured against an ordinary male non-dynastic leader. This omitted group is the largest of the four, as Section 6.3.1 shows, so it is the closest thing in the data to an average district head.

In Table 10, both female groups have positive estimates against male non-dynastic leaders while the male dynastic group's estimates cannot be told apart from male non-dynasts. A female dynast is 0.0527 above the omitted group, with a standard error of 0.0184, and a female non-dynast is 0.0370 above it, with a standard error of 0.0150. A male dynast is 0.0159 with a standard error of 0.0139, which is positive but too imprecise to distinguish from zero.

Panel B turns those levels into the two contrasts: gender component vs dynasty component. However, neither of them is precisely estimated, so the results can only be interpreted as an ordering as opposed to a pair of tests. The gender contrast, holding dynastic status fixed, is 0.0368 with a standard error of 0.0196, and the dynasty contrast, holding gender fixed, is 0.0157 with a standard error of 0.0214. The gender contrast is more than twice the size of the dynasty contrast and it is the one closer to conventional significance, at a p-value of 0.061 against 0.463. Their confidence intervals also overlap across most of their range, so I cannot reject that the two are equal, and consequently this table on its own does not establish that the dynastic tie is irrelevant.

Table 10: Leader groups, relative to a male non-dynastic leader

	Static TWFE
Panel A. Leader groups, relative to a male non-dynast	
Female dynast	0.0527***
	(0.0184)
Female non-dynast	0.0370**
	(0.0150)
Male dynast	0.0159
	(0.0139)
Panel B. Contrasts	
Gender, holding dynasty fixed	0.0356*
	(0.0200)
Dynasty, holding gender fixed	0.0145
	(0.0214)
The interaction	−0.0024
	(0.0235)
Observations 3,710 • clusters 514 • $F(3, 513) = 4.07$, $p = 0.007$	

Static two-way fixed effects on the aggregate index, with district and wave fixed effects and districts weighted by village count. The omitted group is elections with no woman and no dynastic descendant on the ticket, so every coefficient in Panel A is measured against an ordinary male non-dynastic leader and none is comparable to the estimates in Table 7. Panel B reports linear combinations of the Panel A coefficients. Standard errors, in parentheses beneath each estimate, are clustered on district. Significance: *** 1%, ** 5%, * 10%.

The table also shows that, on average, the two characteristics do not reinforce each other: the interaction is −0.0002 with a standard error of 0.0234. When interpreted alongside Figure 3, where female dynasts are the only category whose joint test rejects while male dynasts cannot be separated from zero, the two exercises point the same way. On average, the effect is carried more by electing a woman than by the family she comes from.

These averages leave the role of the family open. In the model in Section 2, a family's network adds capacity and enables capture at the same time, so its net effect can go either way, and it helps a district head more when she places more weight on delivering public goods. A small average dynasty contrast is therefore consistent with two different situations. The family may matter little for every district head, or it may help some district heads and not others, so that the average across them is small. The same is true of the interaction. The model predicts that a woman who governs in her own right gains more from her family's network than a man does, while a woman who governs on her family's behalf gains the same as a man. An interaction near zero on average does not tell these situations apart, and which families and which women carry the effect is a question that these averages cannot answer.

6.4 Robustness checks

There are design choices in the main regression specification that might affect the results. I check in three ways how robust the main estimates are when the choices differ, and Appendix Table A3 reports all of them across the four outcomes. The first check drops the village weights. The main specification weights each district by the number of villages it contains, so that every village counts equally. An obvious question is how much of the result is due to the weighting. When the regression is run unweighted, the average cumulative estimate for the aggregate index is 0.0721 against 0.0747 in the main specification, and the standard error roughly doubles, from 0.0175 to 0.0338. The point estimates for the other three outcomes hold as well, with education larger unweighted at 0.0878 and household smaller at 0.0454. This shows that the weighting affects precision and not the result.

The second check I run changes the unit of observation. I rebuilt the outcome panel at the sub-district level, which gives me 7,567 units instead of 514. Unlike the district, the sub-district is a territorially stable unit that district proliferation does not reshape. The panel starts in 2006 because the earlier waves cannot be linked forward reliably. The average cumulative estimate for the aggregate index when the unit is at the sub-district level is 0.0611 with a standard error of 0.0181, against 0.0747 at district level. Health and household are close to their district values, at 0.0219 and 0.0378 (versus 0.0222 and 0.0728). Education is the one that moves, falling to 0.0275 against 0.0574 at the district level. This is because a third of sub-districts have no education value in the 2008 wave, which is the wave the district-level education estimate leans on most heavily. I estimated the sub-district panel on an earlier vintage of the treatment file, so its estimates are not refreshed to the ones in Table 7.

The third check widens the treatment definition. The main definition treats a district when it elects a female dynast, which excludes women who are the first in their family to hold office. Adding the founders whose relative had already entered office by the time of their own election, which brings in three more districts, gives an aggregate estimate of 0.0726 against 0.0747. Education is 0.0623, health is 0.0244, and household is 0.0576. Changing the treatment definition does not change sign, significance, or magnitude by more than a few percent.

6.5 Pre-trends

The estimator reports three placebo coefficients before the election, and they are not all small. On the aggregate index the lead at $t-3$ is -0.0613 with a t of -1.71, and on household

infrastructure it is -0.1199 with a t of -1.75. In both cases that is the largest coefficient on the pre-election side of the panel. The joint placebo tests do not reject for any of the four outcomes, at 0.227, 0.526, 0.678 and 0.279. However, a joint test can pass while one well-supported placebo is large, and that is why the event studies in Figure 1 show every placebo coefficient next to the joint tests in Table 7.

I run an alternative estimate for the pre-trend using a static two-way fixed effects event study in which the weights enter the regression directly. This specification uses only the pre-switch observations of districts that eventually elect a female dynast, together with every wave of the districts that never do, which means that no observations after an election can contaminate the leads. The three leads are -0.0328, -0.0256 and 0.0002, none of them individually significant, and the joint test that all three are zero does not reject at $p = 0.282$. Estimated without weights the leads are -0.0385, -0.0220 and -0.0018, and the joint test gives $p = 0.515$. The second estimator therefore does not find a pre-trend that the first one missed, and it does not find one that can be distinguished from zero.

A key takeaway from this exercise is that the pre-period is not clean enough to set aside, but also not strong enough to overturn the results. The lead at $t - 3$ is negative in both of the outcomes where it is largest, it reflects something real about the districts that elect a female dynast, the effects I report are understated. I cannot rule out that it reflects a trend that would have continued into the post-election period on its own. Settling that needs a design that does not rest on parallel trends. The dataset records every candidate pair that contested each election and the share of the vote it received, so a close-election regression discontinuity is feasible in principle, and that is where I would take this next.

A takeaway from this exercise is that the pre-period is not clean enough to set aside *and* not strong enough to overturn the results. The largest of the three leads for aggregate index and household infrastructure is at $t - 3$, with the coefficients equal to -0.0613 ($t = 1.71$) and -0.1199 ($t = 1.75$), respectively. A negative lead at $t - 3$ means that relative to the comparison districts, districts that elect a female dynast had lower outcomes three rounds before the election than in the round immediately preceding it. Thus, their outcomes were therefore already rising before the election. If that rise reflects a trend that would have continued in the absence of the election, part of the estimated post-election effect is attributable to the trend, and the reported effects are overstated. Resolving this requires a design that does not rest on parallel trends. The dataset records every candidate pair that contested each election and the share of the vote it received, so a close-election regression discontinuity is feasible in principle and is the natural next step.

7 Heterogeneity analysis

Female dynasts may have different effects on the provision of public goods depending on the context and setting of the district they govern. For instance, a more urbanized district may have a higher baseline of public infrastructure, and the effect of female dynasts may then be more muted there than in less urbanized districts. The effects may also differ between districts in Java and districts outside Java. Java is the island where the capital city is located, and it may be on a different development trend that affects both the election of female dynasts and their effect on the provision of public goods. These heterogeneous effects may show up in one outcome and not in another, and the results can inform us about which public goods constituents in different settings may gain from the election of a female dynast.

In addition to the urbanization of the district and whether it is in Java, I also explore whether the effect differs with the position the female dynast holds on the winning ticket. In Section 3.2, I discuss how the responsibilities and jurisdictions of the head and the deputy, while not identical, are intertwined. For this reason, pooling the head and the deputy in the treatment definition accounts for the potential effect of female dynasts in both positions. The head and the deputy may still have different effects, precisely because their responsibilities and jurisdictions are not identical. The heterogeneity analysis by position can then shed light on the extent to which the gender and dynastic status of the deputy contribute to the overall effect.

In this heterogeneity analysis, I run two types of analysis: a per-arm analysis and a difference analysis. For each split, let $H_g \in \{0, 1\}$ be an indicator for one arm of the split. H_g is defined for every district and does not change over time. In the heterogeneity by urbanization, for instance, $H_g = 1$ if district g is more urbanized and $H_g = 0$ if it is less urbanized.

For the per-arm analysis, I replace the treatment in equation (1) with two arm-specific treatments,

$$D^1_{g,t} \;=\; D_{g,t} \times H_g, \qquad D^0_{g,t} \;=\; D_{g,t} \times (1 - H_g), \tag{9}$$

so that $D^1_{g,t} + D^0_{g,t} = D_{g,t}$ for every district and wave. I then estimate one regression with $D^1_{g,t}$ as the treatment and a second regression with $D^0_{g,t}$ as the treatment. This produces the estimate for each arm on the same scale as the main result. A female-dynast district in the other arm has an arm-specific treatment of zero in every wave, so $F_g = \infty$ in equation (3) and it enters that regression only as a control.

For the difference analysis, I keep the main treatment $D_{g,t}$ and use the district-level estimates that the robust estimator produces for every switching district. At each horizon ℓ,

I regress the switcher-level estimate $\text{DID}_{g,\ell}$ from equation (5) on the arm indicator,

$$\text{DID}_{g,\ell} \;=\; \kappa_\ell \;+\; \theta_\ell H_g \;+\; u_{g,\ell}, \qquad g \in \mathcal{S}_\ell, \quad \ell = 1, \ldots, 4. \tag{10}$$

The coefficient θ_ℓ is the difference between the two arms in the effect at horizon ℓ, and it comes with its own standard error.[3] I run the same regression on the switcher-level placebos, where the dependent variable is the term in square brackets in equation (7) for $\ell = 1, 2, 3$. A nonzero coefficient there means that the two arms were already on different paths before the election. Reading θ_ℓ as a difference in effects requires that the outcome evolution that switching districts would have had without a female dynast is uncorrelated with H_g, and the placebo regressions are a test of that assumption. With the difference analysis, I can show whether the estimates of the two arms are statistically different from each other, and whether the split also predicts a difference before the election. The arm estimations and the difference analysis both use the main specification in Section 5.1 with the robust estimator of De Chaisemartin and d'Haultfoeuille (2020).

7.1 Heterogeneity by urbanization

Statistics Indonesia classifies every village as urban or rural with a score built from three indicators: population density, the percentage of households working in agriculture, and access to urban facilities (Badan Pusat Statistik, 2021*b*). A village receives between one and eight points for its population density and between one and eight points for the share of its households outside agriculture. It then receives one point for each of nine urban facilities. Six of these are facilities within a set distance, which are a kindergarten, a junior secondary school, a senior secondary school, a market, shops, and a hospital. The other three are the presence of a hotel or entertainment venue, at least two percent of households with a fixed-line telephone, and at least 95 percent of households using electricity from the state electricity company. A village with a total score of nine or more is urban, and a village with a lower score is rural. The classification therefore summarizes how dense, how non-agricultural, and how well served a village is. Each of these can affect the selection of female dynasts and the provision of public goods differently across districts. In turn, it may affect how female dynasts behave when governing, and ultimately the kind of effects they bring to districts with different levels of urbanization.

[3] This is the `predict_het` option of the `did_multiplegt_dyn` command. The regression runs over switching districts only, so the number of districts behind it falls as the horizon lengthens. The dependent variable is itself estimated, and it is not established whether the reported standard errors account for this. If they do not, the inference on the difference is optimistic.

Indonesian districts generally contain both urban and rural villages. To conduct the heterogeneity analysis by urbanization, I therefore calculate the share of villages in each district that Statistics Indonesia classifies as urban in the 2021 PODES wave. I assign a district as "more urbanized" if its urban share is above the median share across all 514 districts, which is 0.118, and as "less urbanized" otherwise. Table 11 summarizes how urbanization looks across Indonesian villages and districts. Of the 84,096 villages in the 2021 wave, 16,359 are urban, which is 19.5 percent. The two groups of districts differ sharply. In more urbanized districts, the mean urban share is 0.489. In less urbanized districts, it is 0.049. Of the 76 districts that elect a female dynast, 50 are more urbanized and 26 are less urbanized. All 15 treated *kotamadya* are in the more urbanized group.

Table 11: Urbanization of villages and districts

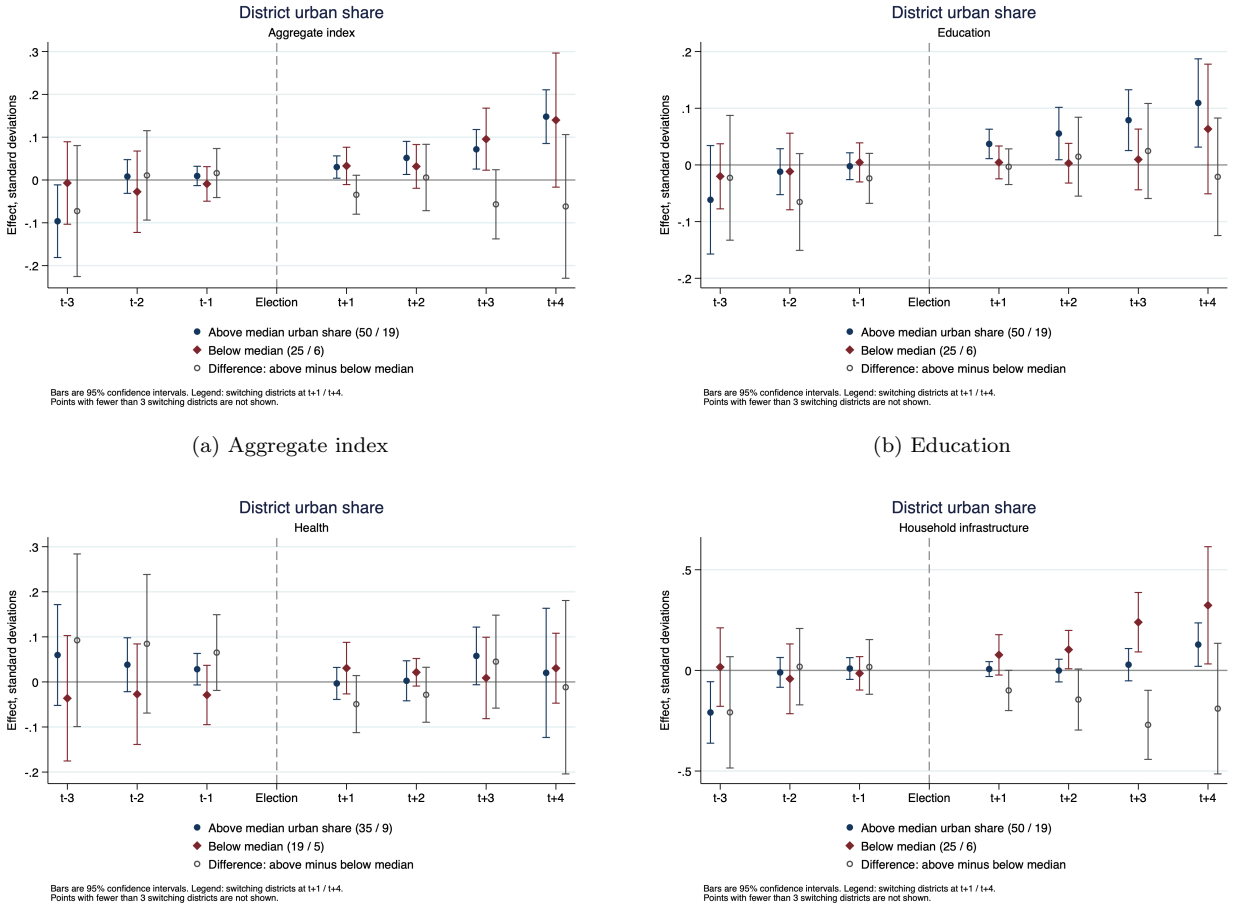
	All districts		Treated districts	
	More urbanized	Less urbanized	More urbanized	Less urbanized
Districts	257	257	50	26
Villages	39,423	44,673	9,813	3,816
Urban villages	14,125	2,234	3,797	224
Mean urban share	0.489	0.049	0.505	0.057
Kotamadya	96	2	15	0
Household count, 2011 (0–5)	3.51	2.14		
Education index, 2011	0.446	−0.372		

Villages are classified with the urban code in the 2021 PODES wave. A district is more urbanized if its share of urban villages is above the median share across all 514 districts (0.118). Treated districts are the districts that elect a female dynast in the sample period. The last two rows are district means in the 2011 wave: the household count is the number of the five household amenities in Section 4.6, and the education index is in village standard deviations.

The effect of female dynasts is approximately the same in more and less urbanized districts for the aggregate index. The average effect on this outcome is 0.074 (0.020) in more urbanized districts and 0.065 (0.030) in less urbanized districts. The components of the index do not move together, however. The effect on household infrastructure is concentrated in less urbanized districts, at 0.167 (0.061) against 0.031 (0.032) in more urbanized districts, and it grows with time in office: the cumulative effect in less urbanized districts is 0.077, 0.103, 0.239 and 0.323 at the four lags. The effect on education goes in the opposite direction. It is 0.074 (0.025) in more urbanized districts and 0.012 (0.023) in less urbanized districts. The effect on health is small in both arms, at 0.020 and 0.024. Figure 4 shows the event study for each arm, together with the difference test.

The difference test separates the two arms only on household infrastructure. The differ-

ence between more and less urbanized districts is negative at all four lags, at −0.099, −0.145, −0.270 and −0.190. The third lag is significant at the one percent level, and the first two are significant at the ten percent level. The fourth lag rests on 25 districts and is not significant. On the aggregate index and on education, the difference test finds no significant difference at any lag.



(a) Aggregate index

(b) Education

(c) Health

(d) Household infrastructure

Figure 4: Heterogeneity by urbanization

Each colored series is a separate estimation in which the treatment is a district that elects a female dynast and belongs to that arm, with every other district as the control. Each point is the cumulative effect at that horizon, and negative horizons are leads. The grey series with hollow markers is the difference test: the coefficient on an indicator for the more urbanized arm in a regression of the district-level effect estimates on that indicator, run over the districts that elect a female dynast. Bars are 95 percent confidence intervals, clustered on district. The legend gives the number of switching districts at the first and fourth lag. Points with fewer than three switching districts are not shown.

The results suggest that female dynasts cause significantly larger gains in household infrastructure in less urbanized districts. One reason is the starting point. In 2011, the average less urbanized district had 2.14 of the five household amenities, against 3.51 in the average

more urbanized district, so less urbanized districts had more room to add them. The urban classification itself does not use any of the household amenities. Instead it uses the presence of schools, which are what the education index counts, and the education effect is the one that is larger in more urbanized districts. Because household infrastructure and education move in opposite directions across the two arms, it thus follow that the ordering does not show up in the aggregate index.

For household infrastructure, the difference analysis gives a negative coefficient at every lag, which means the effect is larger in less urbanized districts throughout. The coefficients are -0.099, -0.145, -0.270 and -0.190 at $t+1$ through $t+4$. The first two are significant at the ten percent level, the third is significant at the one percent level, and the fourth is not significant.

The arm estimates for education differ, but the difference analysis does not confirm the gap. Its coefficients are close to zero at every lag. The two analyses compare different things. Each arm is compared with all districts that do not switch into it, so the control group differs between the two arms. The difference analysis compares female-dynast districts with each other. The source of the disagreement on education is not established, and for this reason I do not interpret the education ordering as a difference in the effect of female dynasts.

However, there is one caution. In more urbanized districts, the household-infrastructure estimate has a lead of -0.209 at $t-3$ with a standard error of 0.078, which rests on 11 districts. The negative lead means that these districts were already improving relative to the comparison districts before the election, so the effect in this arm may be overstated. The difference analysis has a lead of -0.208 at the same horizon, which is not significant, and it means that more urbanized districts were improving faster than less urbanized ones before the election. Both leads therefore point toward the gap in household infrastructure being understated.

7.2 Heterogeneity by Java and outside Java

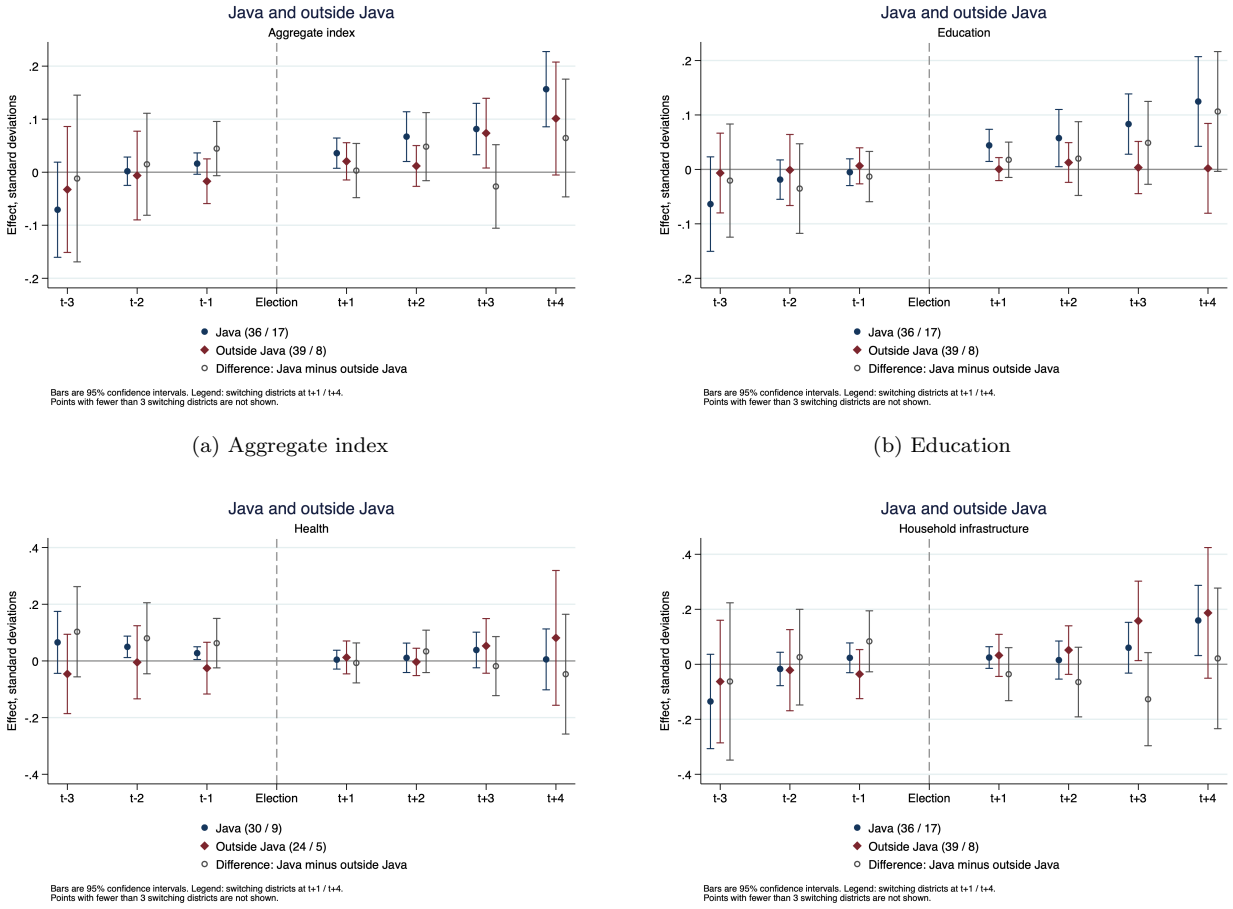
Java, the island where the capital city of Jakarta is located, has been the center of Indonesia's development since before independence. The Dutch East India Company captured the port town of Jayakarta in 1619 and rebuilt it as Batavia, the headquarters of its trade in Asia, and Dutch colonial rule expanded outward from Java to the rest of the archipelago (Ricklefs, 2008). Colonial investment on Java also left infrastructure and industry that persist today (Dell and Olken, 2020). It holds 56.1 percent of the population (Badan Pusat Statistik, 2021*a*), and it produces about 59 percent of national output (Badan Pusat Statistik, 2021*c*). Manufacturing is concentrated there, and Deichmann et al. (2005) find that improvements

in transport infrastructure do little to draw industry to centers outside Java. The same gap shows up in the PODES data. In 2011, the average Java district had 3.47 of the five household amenities, against 2.63 outside Java, and its education index was 0.867 village standard deviations, against −0.213 outside Java. Java is also highly urbanized. Of the 119 districts in Java, 113 are more urbanized by the definition in Section 7.1, against 144 of the 395 districts outside Java. If urbanization alone drives the heterogeneity, Java should look like the more urbanized arm, with a larger education effect and a smaller household-infrastructure effect than districts outside Java.

The Java split may still differ from the urbanization split, for three reasons. The first is composition. Of the 36 treated districts in Java, 34 are more urbanized, while the 40 treated districts outside Java include 16 more urbanized and 24 less urbanized districts. The Java arm is therefore close to a subset of the more urbanized arm, and the arm outside Java mixes both kinds of district. The second is proximity to the capital, which may give district governments in Java easier access to central ministries and to national funding. The third is the historical position of Java as the center of political power, which may shape how established its political families are (Cirone and Pepinsky, 2022). The data here can test the first reason. The second and third are possible explanations that this analysis does not test.

The effect of female dynasts on the aggregate index is larger in Java, at 0.089 (0.023), than outside Java, at 0.039 (0.023). The education effect is concentrated in Java, at 0.083 (0.028) against 0.006 (0.018) outside Java. The household-infrastructure effect is larger outside Java, at 0.087 (0.051), than in Java, at 0.060 (0.040), and the effect on health is small in both arms, at 0.018 and 0.025. The Java arm fails its joint placebo test on the aggregate index ($p = 0.013$) and on health ($p = 0.017$). The difference test does not separate the two arms on any outcome. On the aggregate index, the difference between Java and outside Java is 0.003, 0.048, −0.027 and 0.064 at the four lags, and none is significant. On education, it is 0.018, 0.020, 0.049 and 0.106, and only the fourth lag is significant, at the ten percent level. On household infrastructure, it is −0.036, −0.065, −0.127 and 0.021, and none is significant. Figure 5 shows the results.

The Java results go in the direction that urbanization predicts. Education gains are concentrated in Java, and household-infrastructure gains are larger outside Java. The household gap between the two arms is much smaller than the gap between the urbanization arms, however, at 0.087 against 0.060 compared with 0.167 against 0.031. The composition of the arms accounts for this. Most of the less urbanized treated districts are outside Java, but the arm outside Java also contains 16 more urbanized districts, and these dilute the household effect of that arm. Since almost every treated district in Java is more urbanized, this analysis



(a) Aggregate index

(b) Education

(c) Health

(d) Household infrastructure

Figure 5: Heterogeneity by Java and outside Java

Java is provinces 31 to 36, including Banten. Each colored series is a separate estimation in which the treatment is a district that elects a female dynast and belongs to that arm, with every other district as the control. Each point is the cumulative effect at that horizon, and negative horizons are leads. The grey series with hollow markers is the difference test: the coefficient on an indicator for Java in a regression of the district-level effect estimates on that indicator, run over the districts that elect a female dynast. Bars are 95 percent confidence intervals, clustered on district. The legend gives the number of switching districts at the first and fourth lag. Points with fewer than three switching districts are not shown.

cannot separate a Java effect on education from an urbanization effect. The larger aggregate effect in Java should also be read with the failed placebo test of the Java arm in mind. Taken together, the Java split does not add a separate dimension of heterogeneity beyond what urbanization shows.

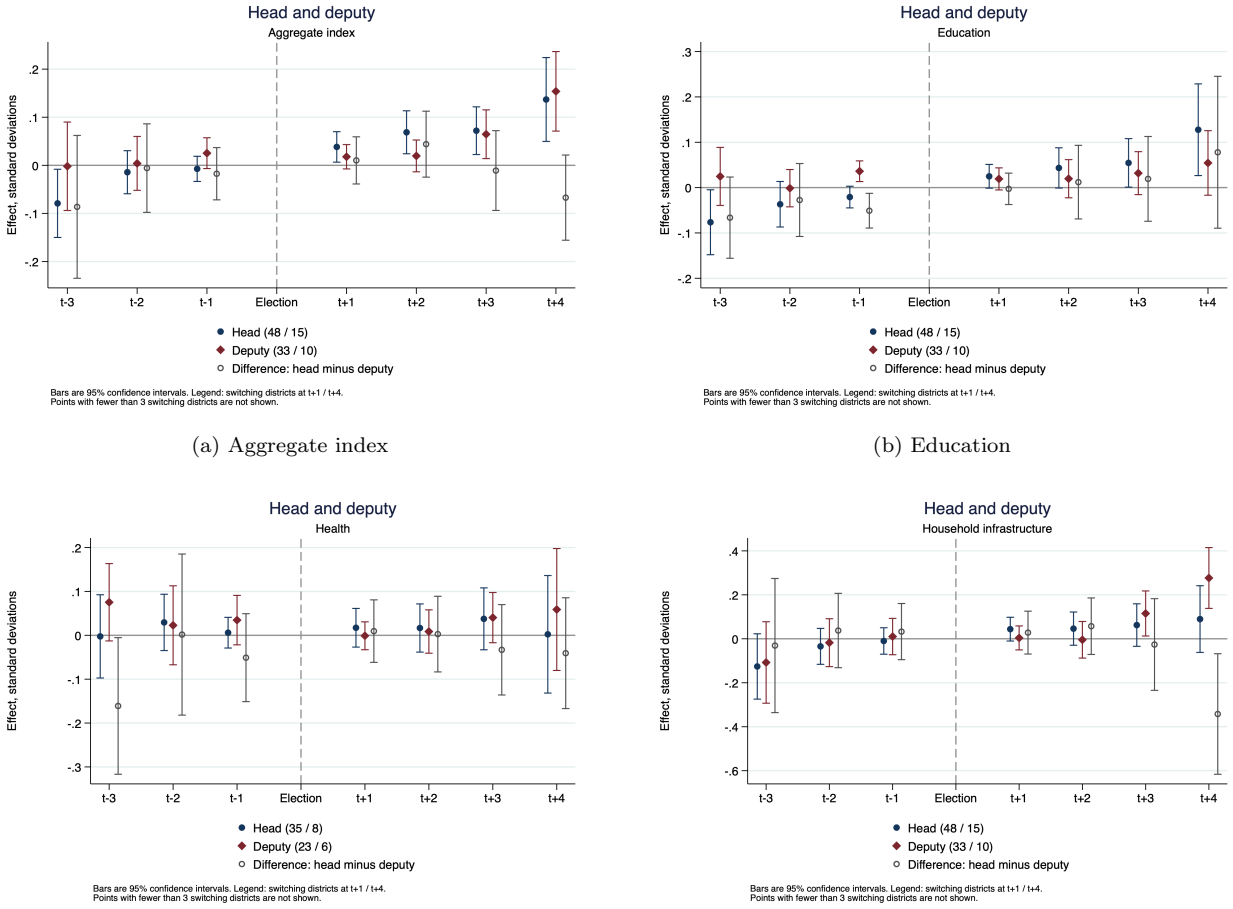
7.3 Heterogeneity by position on the ticket

The treatment definition pools female dynasts elected as head and as deputy, and Section 3.2 explains why. Looking at the two positions separately is still informative, because the law

gives them different powers. Under Article 65 of Law 23 of 2014, the head proposes the development plans, the annual work plan and the budget. Under Article 66, the deputy assists the head, coordinates the district agencies, and monitors and evaluates the administration of the district agencies and the villages. If the effect of female dynasts works through the powers of the office, the effect should be larger when the female dynast is the head. If it works through something that the head and the deputy share, such as the political family behind the ticket, the two effects should be similar. The timing of the effect may also differ, since the deputy's influence works through monitoring and coordination and the head's works through the budget. Of the 76 treated districts, 49 elect a female dynast as head at least once and 33 elect one as deputy at least once. Six districts elect a female dynast to each position in different election cycles, and they appear in both arms. The difference test drops these six districts so that the position is fixed within each district.

The effect on the aggregate index is 0.079 (0.024) when the female dynast is the head and 0.061 (0.021) when she is the deputy. The education effect is 0.058 (0.025) for heads and 0.036 (0.024) for deputies, but both arms fail their joint placebo tests on education ($p = 0.029$ and $p = 0.017$). The household-infrastructure effect is 0.064 (0.043) for heads and 0.080 (0.045) for deputies, and the health effect is 0.024 for heads and 0.023 for deputies. The two arms also differ in timing. On the aggregate index, the deputy arm is 0.018 and 0.020 at the first two lags, against 0.038 and 0.069 for the head arm, and it then rises to 0.065 and 0.154 at the third and fourth lags. On household infrastructure, the deputy arm is close to zero at the first two lags and reaches 0.115 and 0.277 at the third and fourth. The difference test does not separate the two positions. On the aggregate index, the difference between head and deputy is 0.010, 0.044, -0.011 and -0.067, and none is significant. On household infrastructure, it is 0.028, 0.057, -0.026 and -0.342, and only the fourth lag is significant, at the five percent level, on 21 districts. Figure 6 shows the results.

The results do not show that female dynasts elected as head have a larger effect than those elected as deputy. The arm estimates are close, and the difference test finds no consistent difference across lags. The one significant difference, at the fourth lag on household infrastructure, rests on 21 districts and has the opposite sign from the first two lags. This supports the pooling of the two positions in the treatment definition. If the effect worked mainly through the budget powers of the head, the head arm should be clearly larger. The timing differs, however. The deputy effect arrives later, at the third and fourth lags. One possible reading is that a deputy's influence through monitoring and coordination takes longer to show up in the stock of public goods than the head's influence through the budget.



(a) Aggregate index

(b) Education

(c) Health

(d) Household infrastructure

Figure 6: Heterogeneity by position on the ticket

Each colored series is a separate estimation in which the treatment is a district that elects a female dynast to that position, with every other district as the control. Six districts that elect a female dynast to each position in different cycles appear in both arms. The grey series with hollow markers is the difference test: the coefficient on an indicator for the head position in a regression of the district-level effect estimates on that indicator, run over the districts that elect a female dynast, with the six districts dropped. Each point is the cumulative effect at that horizon, and negative horizons are leads. Bars are 95 percent confidence intervals, clustered on district. The legend gives the number of switching districts at the first and fourth lag. Points with fewer than three switching districts are not shown.

8 Mechanisms

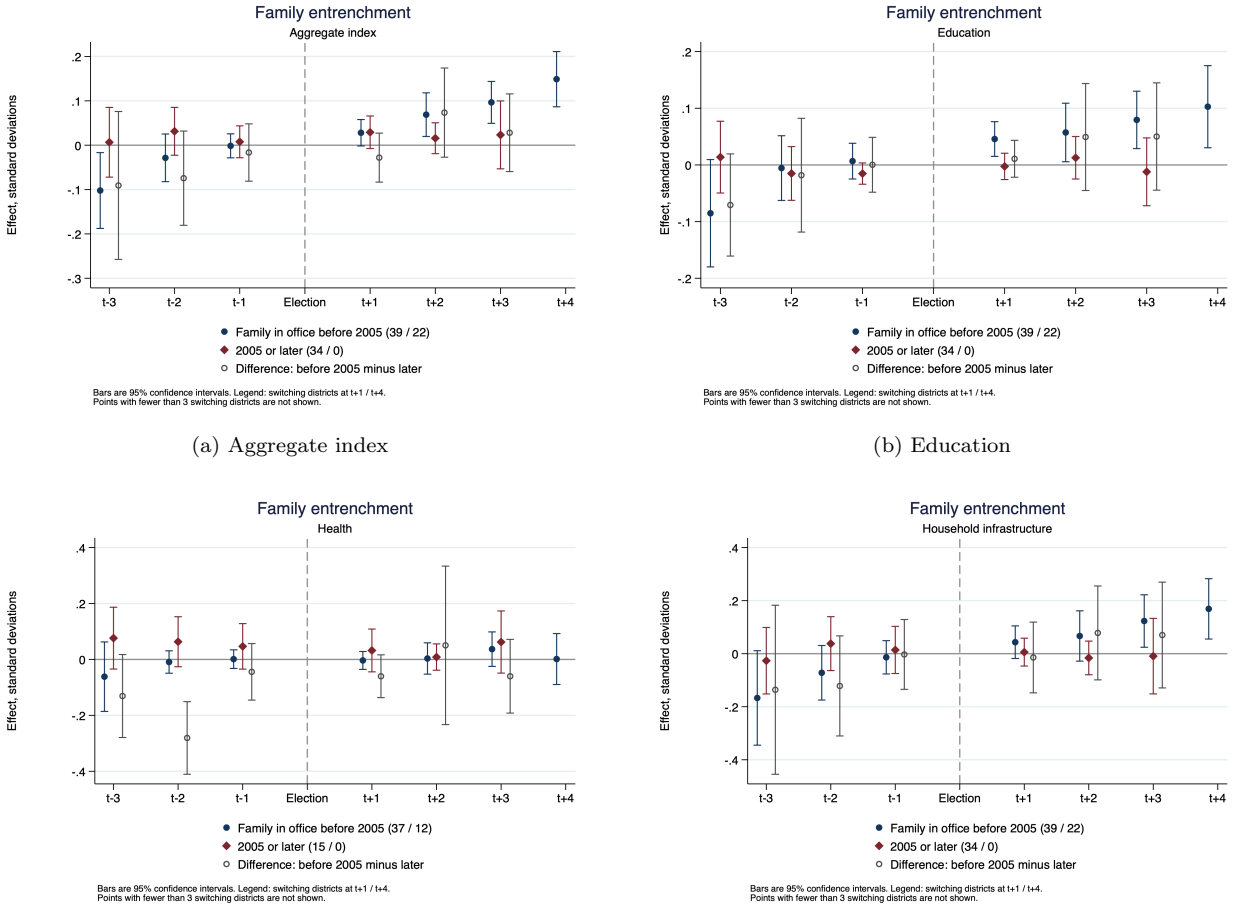
The model in Section 2 leaves one question open while making two predictions. Because the network adds capacity and enables capture at the same time, whether a family's network helps or hurts the provision of public goods is left ambiguous. The model then predicts that a family's network raises the public goods a district head delivers by a larger amount when she places more weight on delivering them, because a district head who cares more about delivery diverts less of the budget to rents. It also predicts that a woman who governs in

her own right gains more from her family's network than a man does, while a woman who governs on her family's behalf gains the same as a man. In this section, I first ask what kind of family helps the provision of public goods. I then ask whether a family's network helps women more than it helps men, and whether the gain comes from women who govern in their own right. Lastly, I describe the women who carry the effect and examine two alternative channels.

8.1 Which dynasties? Capacity against capture

The first result of the model says that the net effect of a family's network is an empirical question. As discussed in Section 2, the forces within the family network are strongest in families that have held office for a long time due to the time they had accumulating power and network. It follows that the net effect of the capacity and capture forces should be most visible on families that held office prior to the direct election. To show this, I classify a female dynast's family as entrenched if it first held public office before 2005, using the year in which the family first held office as recorded in the dynasty data, and as newer if it first held office in 2005 or later. When a district elects a female dynast more than once, I classify the district as entrenched if any of those families first held office before 2005. I estimate the effect separately for each group, with that group as the treatment and every other district as the control, and I run the difference analysis from Section 7 with the entrenchment indicator. Figure 7 shows the results.

Figure 7 shows that the effect is concentrated in districts where the female dynast's family is entrenched. On the aggregate index, the entrenched group rises at every lag, from 0.028 at the first lag to 0.149 at the fourth, and the confidence intervals exclude zero from the second lag onward. The effect on newer family stays between 0.016 and 0.029 through the third lag, with confidence intervals that contain zero. The average cumulative effect per treated district, which averages across the lags and is not shown in the figure, is 0.0942 with a standard error of 0.0239 in the entrenched group, against 0.0258 (0.0220) in the other group. Education and household infrastructure follow the same pattern. In the entrenched group, education rises from 0.046 to 0.103 and household infrastructure from 0.043 to 0.169, while in the other group neither moves away from zero. Health is small in both groups, and its confidence intervals contain zero throughout. This answer the open question from the model: for female dynasts, the capacity force outweighs the capture force.



(a) Aggregate index

(b) Education

(c) Health

(d) Household infrastructure

Figure 7: Effect by the female dynast's elite entrenchment timing

Each colored series is a separate estimation in which the treatment is a district that elects a female dynast whose family first held office in that period, with every other district as the control. Each point is the cumulative effect at that horizon, and negative horizons are leads. The grey series with hollow markers is the difference analysis: the coefficient on an indicator for a family that first held office before 2005 in a regression of the district-level effect estimates on that indicator, run over the districts that elect a female dynast. Bars are 95 percent confidence intervals, clustered on district. The legend gives the number of switching districts at the first and fourth lag. Points with fewer than three switching districts are not shown.

8.1.1 Breadth and depth of the family network

If old families help a female dynast through their capacity, the next question is what that capacity consists of. A family's network can be deep, in the sense that the family has held the very seat she now holds and knows its bureaucracy and its budget process. It can also be broad, in the sense that the family's members hold office in legislatures, in other districts, and at the provincial and national levels. Table 12 describes the families of the female dynasts in the two groups.

Entrenched families are broader than newer families, but they are not deeper. On breadth,

Table 12: Breadth and depth of the family network, by entrenchment

	Entrenched family	Newer family
Breadth		
Office-holding kin	3.51	1.74
of whom in executive office	1.95	1.32
of whom in legislative office	2.07	0.50
Any kin outside the district	0.44	0.26
Kin at provincial level	0.29	0.11
Kin at national level	0.05	0.00
Depth		
Family held this district's executive office before her	0.73	0.74
Family held it in the term immediately before her	0.36	0.51
Elections	55	40

Means per treated female-dynast election, not estimates. Entrenched: the female dynast's family first held public office before 2005; newer: in 2005 or later. The breadth rows count the female dynast's office-holding relatives, following family links outward (Section 4.2); each is a lower bound, because the records give a relative's first office and not her whole career, and two elections in the newer group have no breadth record. The depth rows are shares: whether a member of her family held the executive office of the district she now leads, at any time before her or in the term immediately before her.

a female dynast from an entrenched family has an average of 3.51 office-holding relatives, against 1.74 for a female dynast from a newer family. Most of the difference is in legislatures, at 2.07 relatives against 0.50, and entrenched families also reach further outside the district and up to the provincial level. Part of this gap is mechanical, because a family that entered politics earlier has had more years to place its members in office. On depth, the two kinds of family look alike. Among female dynasts from entrenched families and among those from newer families alike, about three quarters come from a family that had already held the executive office of the district she now leads. A female dynast from an entrenched family is in fact less likely to have taken the seat directly from a relative, at 36 percent of elections against 51 percent for a female dynast from a newer family. What sets entrenched families apart is therefore the breadth of their network.

If the breadth of the network is what an old family passes on that makes a female dynast provide more public goods, the depth of the network can serve as a falsification test. A female dynast whose family has held her seat, and in particular one who took the seat directly from a relative, should then have a smaller effect on the provision of public goods. To check this, I sort the 102 treated elections by the family's prior hold on the district's executive office: her family held the office in the term immediately before her (direct handover, 41 elections), her family held it in the past and at least one term passed between that relative and her

(delayed handover, 31 elections), or she is the first in her family to hold it (no prior claim, 29 elections). The effect on the aggregate index is largest in districts with no prior claim, at 0.118 (0.038), against 0.077 (0.029) for a direct handover and 0.045 (0.025) for a delayed handover (Appendix Table A8). The ordering holds under the alternative specifications in Appendix Table A7. The women whose families never held the seat also come from the broadest networks, with 4.38 office-holding relatives against 2.20 for the direct-handover group (Appendix Table A9). Appendix G also shows a check on how breadth works through the district budget: districts that elect a woman whose family never held the seat shift spending toward capital.

8.2 The family or the woman?

The second result of the model says that a family's network raises the public goods a district head delivers by more when she places more weight on delivering them. A district head who places more weight on delivery diverts less of the budget to rents, so more of the capacity that her family brings reaches her constituents. If a woman places more weight on delivery than a man, an old family should therefore help a female dynast deliver more public goods than it helps a male dynast. The specificity test in Section 6.3 cannot tell whether this is the case. That test separates the effect of being a woman from the effect of being a dynast, and it finds that on average the two do not reinforce each other. The average pools old families with newer ones, so it does not show whether an old family helps women more than it helps men.

To test whether an old family helps women more than men, I extend the joint regression from Section 6.3 by splitting female dynasts and male dynasts each by whether the family is entrenched. I keep female non-dynasts as one group because entrenchment is defined only for dynasts, and male non-dynasts are omitted so every coefficient is measured against an ordinary male non-dynastic district head. Table 13 reports the results.

Similar to the specificity test, I estimate this regression with static two-way fixed effects. There are two reasons for this choice. First, the comparison needs all four dynast groups in one regression, and De Chaisemartin and d'Haultfoeuille (2020)'s heterogeneity-robust estimator takes a single treatment. Second, splitting the treated districts by gender and by entrenchment thins the support. The heterogeneity-robust estimator identifies the effect at each horizon from the districts that switch into treatment at that horizon, and it identifies each placebo from the districts observed before they switch. Once dynasts are split by gender and by entrenchment, the cells hold between 37 and 106 districts, and the number of districts observed at the later lags and the earlier leads is a fraction of that. In cells this small, the

horizon-by-horizon estimates and their placebos rest on a handful of districts. Static two-way fixed effects pool the post-election waves into one coefficient for each group, so each coefficient uses all the districts in its cell. With staggered election timing, these coefficients can mix the effects of different horizons and cohorts, so when the effect grows over time, as it does in the main results, a coefficient can understate or overstate the average effect of its group. For this reason, I use these estimates to compare the groups with one another, and I do not compare their magnitudes with the main results.

Table 13: Effect by gender and family entrenchment, relative to a male non-dynastic leader

	Aggregate	Education	Health	Household
Female dynast, entrenched family	0.067***	0.072**	0.021	0.085**
	(0.025)	(0.035)	(0.028)	(0.042)
Female dynast, newer family	0.035	0.024	0.013	0.018
	(0.028)	(0.023)	(0.039)	(0.048)
Male dynast, entrenched family	0.009	0.037**	0.017	-0.036
	(0.016)	(0.018)	(0.016)	(0.027)
Male dynast, newer family	0.020	0.018	0.035	0.016
	(0.021)	(0.022)	(0.025)	(0.037)
Female non-dynast	0.037**	0.040**	0.031**	0.016
	(0.015)	(0.017)	(0.016)	(0.028)
Contrasts				
Entrenched − newer, women	0.032	0.048	0.009	0.068
	(0.036)	(0.040)	(0.046)	(0.062)
Entrenched − newer, men	-0.011	0.019	-0.018	-0.052
	(0.024)	(0.027)	(0.028)	(0.041)
Difference-in-differences	0.043	0.029	0.027	0.120
	(0.041)	(0.044)	(0.051)	(0.074)
Women − men, entrenched	0.058**	0.035	0.004	0.122***
	(0.025)	(0.034)	(0.026)	(0.047)
Women − men, newer	0.015	0.006	-0.022	0.002
	(0.033)	(0.030)	(0.046)	(0.057)
Female dynast (newer) − female non-dynast	-0.002	-0.016	-0.018	0.002
	(0.029)	(0.028)	(0.040)	(0.052)

Static two-way fixed effects, with district and wave fixed effects and districts weighted by village count. The omitted group is elections with no woman and no dynastic descendant on the ticket. Entrenched: the family of the dynast on the ticket first held public office before 2005; newer: in 2005 or later. Elections whose dynast has no usable founding year enter as their own indicators, not shown. Difference-in-differences: the entrenched-minus-newer gap for women minus the same gap for men. Standard errors, in parentheses beneath each estimate, are clustered on district. Significance: *** 1%, ** 5%, * 10%.

Table 13 shows that entrenched families help women provide more public goods but do

not help men in doing so. The first five rows report the effect of each group. On the aggregate index, female dynasts from entrenched families raise provision by 0.067 (0.025), against 0.035 (0.028) for female dynasts from newer families. Male dynasts from entrenched families raise the aggregate index by 0.009 (0.016), and male dynasts from newer families by 0.020 (0.021). The rows under Contrasts compare the groups. The gain from an entrenched family, measured as the effect for entrenched families minus the effect for newer families, is 0.032 (0.036) for women and −0.011 (0.024) for men. The difference-in-differences row subtracts the men's gain from the women's gain, which gives 0.043 (0.041) on the aggregate index and 0.120 (0.074) on household infrastructure. Both are positive. The next two rows subtract the effect of male dynasts from the effect of female dynasts within each kind of family. Within entrenched families, this gives 0.058 (0.025) on the aggregate index and 0.122 (0.047) on household infrastructure, and within newer families it gives 0.015 (0.033) on the aggregate index. The last row shows that female dynasts from newer families are indistinguishable from female non-dynasts, with a difference of −0.002 (0.029).

These results follow the second result of the model. For men, an old family adds nothing on average. The model allows this in two ways: the family may add little capacity and enable little capture, or it may add a great deal of both, so that they offset one another. Meanwhile, for women, the capacity of an old family helps her provide more public goods for her constituents.

8.3 Proxy or own standing?

The third result of the model separates two kinds of female dynast: one who governs in her own right and one who governs as a proxy for her family. A woman who governs in her own right should deliver more than a man from a family with the same network, while a woman who governs on behalf of her family should deliver the same as a man from her family. The case in which governing on the family's behalf is most plausible is direct succession, in which a member of her family held the district's executive office in the term immediately before her.

This case follows the "benchwarmer" logic of Labonne, Parsa and Querubin (2021), who show that term limits bring more women into office in the Philippines, because the female relatives of term-limited incumbents run for the seat to keep it in the family until the incumbent can return. The Indonesian term limit is different. A district head who has served two terms can never hold the same seat again, so there is no return for a female relative to wait for. What carries over to Indonesia is the question of who actually governs during her term: the woman who holds the office, or the relative who held it before her.

Direct succession also allows a test of a different interpretation of the entrenchment result. A female dynast who directly succeeds a relative inherits a governing apparatus that is already in place, such as the officials her relative worked with, the projects he started, and his relationships with the district parliament. If she continues that apparatus, her capacity in the model is higher, and the women who directly succeed a relative should deliver the most public goods. The proxy interpretation predicts the opposite. A woman who directly succeeds a relative is also the woman most likely to govern on his behalf, and in that case she delivers what her family would deliver, which is what a man from her family delivers. The two interpretations therefore make different predictions for the women who directly succeed a relative: the continuation interpretation predicts that they deliver the most, and the proxy interpretation predicts that they deliver the same as men from their families.

I split the entrenched female dynasts and the entrenched male dynasts by whether the ticket directly succeeds a relative, and I estimate the same regression as in Table 13. I use static two-way fixed effects for the same two reasons as before: the comparison needs several groups in one regression, and the cells are small. They are smaller here than in Table 13, because direct succession splits the entrenched dynasts again. The succession indicator is recorded for the ticket, so a female deputy is counted as a direct successor when the head on her ticket succeeds his own relative. Among the 55 entrenched female-dynast elections, 20 directly succeed a relative and 35 do not. Among the 160 entrenched male-dynast elections, 24 directly succeed a relative, 132 do not, and 4 cannot be classified. Table 14 reports the results.

Table 14 shows that the gain from an old family goes through women who did not directly succeed a relative. The first four rows report the effect of each group. On the aggregate index, entrenched female dynasts who did not directly succeed a relative raise provision by 0.090 (0.036), against 0.034 (0.025) for those who did. Entrenched male dynasts are close to zero whether or not they directly succeeded a relative, at 0.008 (0.016) and 0.007 (0.046). The rows under Contrasts compare the groups. Among the tickets that did not directly succeed a relative, subtracting the effect of male dynasts from the effect of female dynasts gives 0.082 (0.037) on the aggregate index and 0.176 (0.055) on household infrastructure. Among the tickets that did directly succeed a relative, the same difference is 0.027 (0.051) on the aggregate index and -0.005 (0.101) on household infrastructure, so women and men who directly succeed a relative are indistinguishable. Within each gender, the women who did not directly succeed a relative deliver 0.056 (0.040) more than the women who did, while for men the same difference is 0.001 (0.045). The difference-in-differences, which subtracts the men's difference from the women's, is 0.054 (0.068) on the aggregate index and 0.181 (0.114)

Table 14: Effect by direct succession and gender, entrenched families

	Aggregate	Education	Health	Household
Female dynast, direct handover	0.034	0.072	-0.012	0.025
	(0.025)	(0.058)	(0.012)	(0.063)
Female dynast, no direct handover	0.090**	0.072**	0.048	0.128***
	(0.036)	(0.037)	(0.045)	(0.049)
Male dynast, direct handover	0.007	0.028	-0.013	0.031
	(0.046)	(0.045)	(0.041)	(0.079)
Male dynast, no direct handover	0.008	0.039**	0.021	-0.048
	(0.016)	(0.018)	(0.016)	(0.030)
Contrasts				
Women − men, no handover	0.082**	0.033	0.027	0.176***
	(0.037)	(0.038)	(0.045)	(0.055)
Women − men, handover	0.027	0.045	0.001	-0.005
	(0.051)	(0.072)	(0.039)	(0.101)
No handover − handover, women	0.056	0.000	0.060	0.102
	(0.040)	(0.063)	(0.041)	(0.074)
No handover − handover, men	0.001	0.011	0.034	-0.079
	(0.045)	(0.043)	(0.040)	(0.085)
Difference-in-differences	0.054	-0.011	0.026	0.181
	(0.068)	(0.089)	(0.072)	(0.114)

Same specification as Table 13, with entrenched female dynasts and entrenched male dynasts each split by whether a relative of the ticket held the district's executive office in the term immediately before. The coefficients on the other groups, which are estimated in the same regression, are not shown. Difference-in-differences: the no-handover-minus-handover gap for women minus the same gap for men. Standard errors, in parentheses beneath each estimate, are clustered on district. Significance: *** 1%, ** 5%, * 10%.

on household infrastructure.

The ordering follows the third result of the model. The women who are the most plausible proxies look like the men from their families, and the women who stand on their own deliver more. The ordering also runs against the continuation interpretation, because the women who directly succeed a relative deliver less than the women who do not.

A second way to identify the women who are most plausibly proxies is to ask whether she is the wife of an officeholder. I record a female dynast as a wife if the relatives listed for her in the dynasty data include a husband. Wives are concentrated in entrenched families, at 24 of the 55 entrenched female-dynast elections against 8 of the 40 newer ones. On the aggregate index, the entrenched female dynasts who are not wives raise provision by 0.093 (0.041), against 0.039 (0.023) for the wives, a difference of −0.054 (0.045). On household infrastructure the two are about equal, at 0.082 and 0.089. The two ways of identifying

proxies overlap only in part: 14 of the 24 entrenched wives directly succeeded a relative, against 6 of the 31 entrenched women who are not wives.

8.4 Who are these women?

The model takes the district head as given and says nothing about who reaches office. There are two ways in which the women who carry the effect could differ from the men from the same families before they take office. The first is that old families may put forward their most qualified women. The second is that voters may hold women to a higher standard, so that the women who win are stronger than the men who win. I examine each in turn.

To compare qualifications, I use the biographical data described in Section 4.2 and Appendix B.6, which cover the winning heads and deputies who are female dynasts, male dynasts, or female non-dynasts. I compare each person at her first win, because a district head who is re-elected looks experienced simply from having held the office. Table 15 reports the means, and Table 16 reports the contrasts.

Table 15: Characteristics at first win, by group

	Women, entrenched, not wife	Women, entrenched, wife	Women, newer family	Men, entrenched	Men, newer family	Female non-dynast
Age	42.6	48.4	46.2	43.3	41.7	46.3
Bachelor's degree or higher	0.71	0.86	0.73	0.87	0.87	0.78
Master's degree or higher	0.50	0.29	0.33	0.34	0.26	0.28
Held public office before	0.50	0.21	0.34	0.54	0.37	0.31
Won a popular vote before	0.50	0.21	0.28	0.47	0.24	0.25
Previously a legislator	0.57	0.21	0.22	0.40	0.26	0.27
Previously a civil servant	0.14	0.07	0.19	0.17	0.23	0.13
Previously in the private sector	0.07	0.14	0.09	0.11	0.26	0.18
Wife of a head or TP-PKK chair	0.07	0.50	0.31	0.06	0.07	0.07
Winning margin (percentage points)	18.1	21.9	13.2	17.4	17.1	14.9

Means over winning heads and deputies in district elections, 2005–2020, at the election at which each person first won. Entrenched: the person's own family first held public office before 2005. Wife: the relatives listed for her in the dynasty data include a husband. Education is the highest level completed before the first win. A characteristic is missing for some people, most often education. The winning margin is the winner's vote share minus the runner-up's.

The entrenched women who are not wives are more experienced than other women. Half of them held public office before, half won a popular vote before, and 57 percent served in a legislature, against 21 percent of the entrenched wives, 22 percent of the women from newer families, and 27 percent of the female non-dynasts. The gap in legislative experience

Table 16: Contrasts in characteristics at first win

	Held public office before	Won a popular vote before	Previously a legislator	Master's or higher
Not wife − wife	0.286*	0.286*	0.357**	0.214
	(0.168)	(0.168)	(0.174)	(0.220)
Not wife − women, newer family	0.156	0.219	0.353**	0.167
	(0.152)	(0.150)	(0.153)	(0.162)
Not wife − female non-dynast	0.195	0.246*	0.305**	0.219
	(0.140)	(0.139)	(0.150)	(0.158)
Not wife − men, entrenched	-0.043	0.033	0.167	0.162
	(0.136)	(0.136)	(0.143)	(0.147)
Men, entrenched − men, newer family	0.177*	0.223***	0.148*	0.080
	(0.091)	(0.084)	(0.085)	(0.098)

Each column is a regression of the characteristic on group indicators without a constant, over the first-win sample. Standard errors, in parentheses beneath each estimate, are clustered on the person. Significance: *** 1%, ** 5%, * 10%.

is 0.36 (0.17) against the wives, 0.35 (0.15) against the women from newer families, and 0.31 (0.15) against the female non-dynasts. Against the men from their own kind of family, they look the same. They held prior office at about the same rate, 0.50 against 0.54. The men from entrenched families are themselves more experienced than the men from newer families, with more prior office, 0.18 (0.09), and more prior popular-vote wins, 0.22 (0.08). Direct succession, unlike wife status, does not separate the women on paper: the entrenched women who directly succeeded a relative and those who did not held prior office at about the same rate, 0.38 against 0.36.

These comparisons say that formal experience does not explain the gap between women and men from old families. The entrenched women who carry the effect bring the same formal experience as the entrenched men, and the men with that experience do not deliver more. In the model, what separates them is the weight a district head places on delivery, which covers both her own preferences and the value she places on building a reputation of her own. The data cannot rule out that these women differ from the men in qualities that a biography does not record.

To examine whether voters hold dynastic women to a higher standard, I use the winner and the runner-up of every district election from 2010 to 2020. Each election contributes two tickets, and I regress whether a ticket won on indicators for the same groups as in Table 13, with election fixed effects, so that each coefficient is the probability that a ticket of that kind beats a male non-dynastic ticket in the same race. Table 17 reports the results.

Dynastic women win as often as dynastic men. Within entrenched families, the difference

Table 17: Winning against a male non-dynastic ticket, top two tickets

	Win	Win, incumbency control	Vote share	Win, heads only
Female dynast, entrenched family	0.439***	0.426***	0.125***	0.325**
	(0.114)	(0.104)	(0.034)	(0.127)
Female dynast, newer family	0.429***	0.415***	0.071***	0.204
	(0.126)	(0.127)	(0.024)	(0.153)
Male dynast, entrenched family	0.462***	0.383***	0.122***	0.412***
	(0.074)	(0.070)	(0.021)	(0.079)
Male dynast, newer family	0.389***	0.382***	0.081***	0.226*
	(0.115)	(0.113)	(0.026)	(0.135)
Female non-dynast	-0.214***	-0.125*	-0.053**	-0.241**
	(0.074)	(0.070)	(0.022)	(0.097)
Contrasts				
Women − men, entrenched	-0.023	0.043	0.003	-0.088
	(0.138)	(0.126)	(0.040)	(0.153)
Women − men, newer	0.041	0.033	-0.010	-0.022
	(0.168)	(0.167)	(0.034)	(0.202)
Difference-in-differences	-0.064	0.010	0.013	-0.066
	(0.220)	(0.214)	(0.054)	(0.256)

Unit: a ticket, the winner and the runner-up of each district election. The first three columns classify the ticket by its head or deputy, 2010–2020 (1,158 elections), because runner-up deputies are coded only from 2010. The last column uses the head only, 2005–2020. Vote share is the ticket's share of the two-ticket vote. Election fixed effects. Tickets whose dynast has no usable founding year enter as their own indicators, not shown. Difference-in-differences: the women-minus-men gap in entrenched families minus the same gap in newer families. Standard errors, in parentheses beneath each estimate, are clustered on district. Significance: *** 1%, ** 5%, * 10%.

between female and male dynasts is −0.023 (0.138), and within newer families it is 0.041 (0.168). The result holds with a control for incumbency, on the vote share, and for heads only. Female non-dynasts are the exception: they lose to male non-dynasts more often, by 0.214 (0.074). There is therefore no sign that voters hold dynastic women to a higher standard than dynastic men, so a voter-imposed standard does not explain why entrenched female dynasts deliver more. Women without a political family do appear to face a higher standard, and a family name appears to remove it. There are two caveats. The comparison conditions on finishing in the top two and on who chose to run. The founding years, and possibly dynastic status, are also coded less thoroughly for runners-up.

8.5 Alternative mechanisms

This section evaluates two alternative mechanisms. I examine whether gender norms are a channel through which the results run, and whether the effect works through alignment between a female dynast's party coalition and the coalition of the provincial governor.

8.5.1 Gender norms

Gender norms can push the effect of electing a female dynast in two directions. On the one hand, where norms about women's roles are more favorable, women may already take part more in public life and in public decisions, so electing one more woman may change less. On the other hand, more favorable norms may reduce the friction that a female district head meets in working with the bureaucracy and the district parliament, which would let her deliver more. In the terms of the model, this second force raises the amount of public goods that she can deliver from each unit of the budget. Gagliarducci and Paserman (2012) show that in Italian municipalities, governments led by a female mayor are more likely to end early, and that this instability is largest where the council is entirely male and in regions with less progressive attitudes toward women's work. Norms may also change who gets elected, because where they are more favorable, voters hold weaker stereotypes against female leaders (Beaman et al., 2009), so a woman needs her family name less to overcome voters' doubts about her gender. Which of these forces dominates is an empirical question.

To measure gender norms, I use the share of women aged 15 to 64 whose main activity in the previous week was work, which I refer to as female employment. I compute it from the individual files of the July rounds of Susenas, the national socioeconomic survey, weighting each woman by her survey weight. Susenas does not ask about attitudes toward women's roles, so I use a behavior that reflects them. Female labor force participation is a common proxy for these norms, and Alesina, Giuliano and Nunn (2013) show that it moves together with attitudes about whether women should work outside the home. I measure female employment before most of the treated elections, and I average it over the 2005, 2006, 2007, and 2008 rounds to reduce sampling noise in a district-level estimate. A district that was later split takes the value of the district it belonged to at the time of the survey. I classify a district as having more favorable norms if its average female employment is above the median across all districts, which is 0.310. Table 18 summarizes the two groups.

I estimate the effect separately for each group, with that group as the treatment and every other district as the control, and I run the difference analysis from Section 7 with an indicator for above-median female employment. Figure 8 shows the results.

On the aggregate index, the effect arrives sooner in districts with higher female employ-

Table 18: Female employment, 2005–2008

	Above median	Below median
Districts	255	256
Mean female employment	0.425	0.236
Range	0.310–0.906	0.094–0.310
Districts that elect a female dynast	38	38
Mean female employment, these districts	0.403	0.223

Female employment is the share of women aged 15 to 64 whose main activity in the previous week was work, from the July rounds of Susenas, weighted by the individual survey weight and averaged over 2005–2008. The median across the 511 districts with a value is 0.310. Three districts have no value.

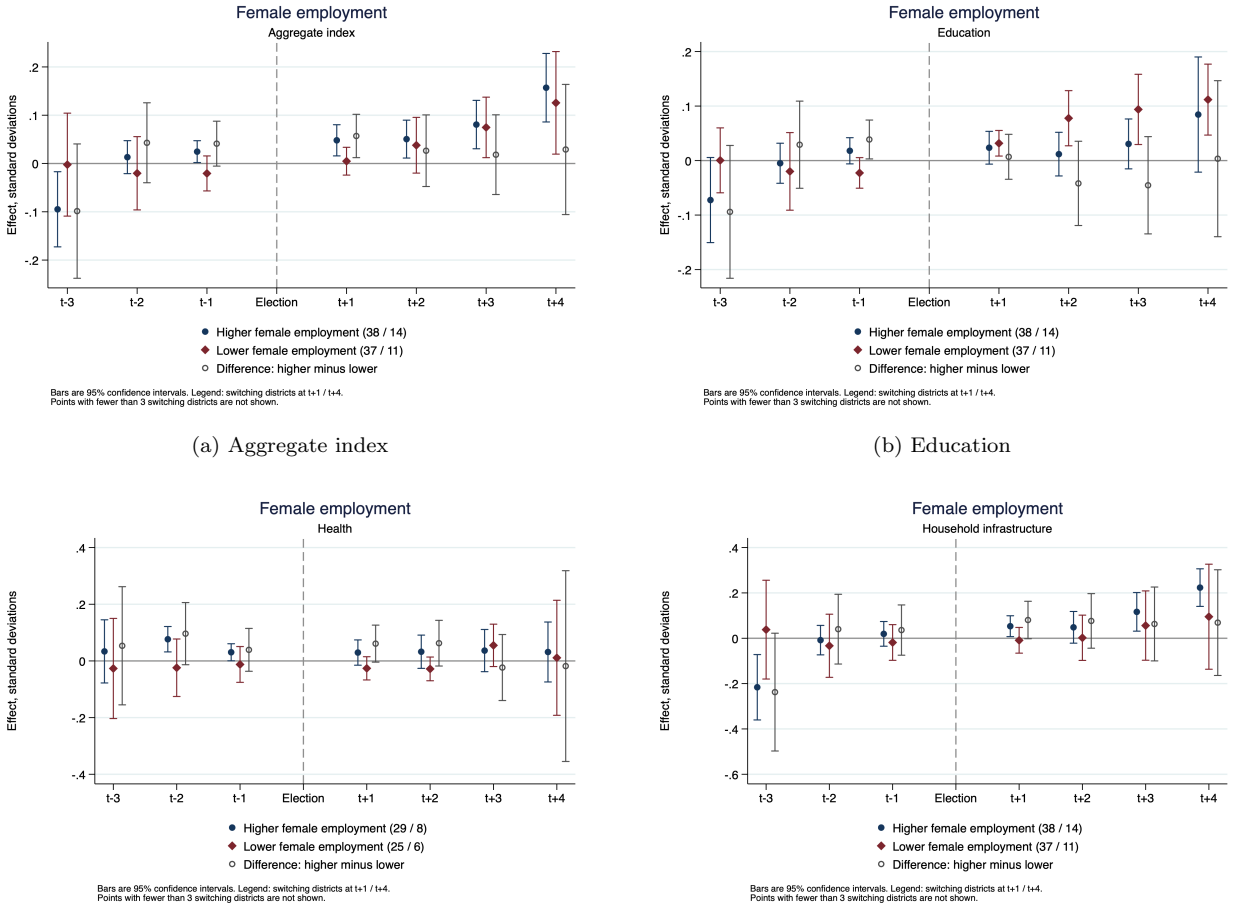
ment, but the two groups end in a similar place. In those districts, the effect is 0.048 at the first lag, and it rises to 0.157 at the fourth. In districts with lower female employment, the effect is 0.005 at the first lag, and it rises to 0.126 at the fourth. The difference analysis shows the same pattern, with a coefficient of 0.057 (0.023) at the first lag and coefficients close to zero at the later lags. The average cumulative effect per treated district, which averages across the lags and is not shown in the figure, is 0.080 (0.021) where female employment is higher and 0.056 (0.029) where it is lower.

The components do not move together. Household infrastructure follows the aggregate index more strongly. It rises from 0.053 at the first lag to 0.223 at the fourth where female employment is higher, while the confidence intervals contain zero at every lag where it is lower. Education goes the opposite way: the effect is larger where female employment is lower, at 0.078 against 0.012 at the second lag. Health is small in both groups.

Taken together, the results suggest that gender norms may affect how quickly a female dynast delivers. They do not show that gender norms change how much she delivers.

8.5.2 Alignment with the provincial government

The second alternative is a political connection that does not depend on the family. Indonesian provinces sit between the central government and the districts, and the governor acts as the representative of the central government in the province. The governor coordinates and supervises the districts, and the province transfers part of its own budget to districts as financial assistance (*bantuan keuangan*). A district head whose winning ticket shares a party with the governor's coalition may find it easier to obtain these funds, to have her programs approved, and to coordinate projects that cross district boundaries. Studies in other countries find that higher levels of government direct more discretionary transfers to local governments that are politically aligned with them, and that aligned constituencies grow



(a) Aggregate index

(b) Education

(c) Health

(d) Household infrastructure

Figure 8: Effect by female employment in the district

Each colored series is a separate estimation in which the treatment is a district that elects a female dynast and has female employment above (or below) the median, with every other district as the control. Each point is the cumulative effect at that horizon, and negative horizons are leads. The grey series with hollow markers is the difference analysis: the coefficient on an indicator for above-median female employment in a regression of the district-level effect estimates on that indicator, run over the districts that elect a female dynast. Bars are 95 percent confidence intervals, clustered on district. The legend gives the number of switching districts at the first and fourth lag. Points with fewer than three switching districts are not shown.

faster (Brollo and Nannicini, 2012; Asher and Novosad, 2017). If alignment is the channel, the effect of electing a female dynast should be larger in districts where her ticket is aligned with the governor.

I code a district election as aligned if the winning ticket shares at least one party with the coalition of the governor in office at the time of the election. For provinces that had already held a direct gubernatorial election, I take the winner of the most recent provincial election held in an earlier year, and I record that ticket's coalition from the pilkada data. I treat a provincial election held in the same year as the district election as not yet in office, because

from 2015 onward the two are held on the same day. For the earliest district elections, held before their province's first direct gubernatorial election, the governor had been elected by the provincial parliament, and I record that governor's coalition from news reports. These elections are almost all in districts that never elect a female dynast: only 2 of the 102 female-dynast elections fall in this period. I leave an election unclassified if the winning ticket or the governor ran as an independent, if the governor was an acting appointee, or if either coalition cannot be identified. Each treated district takes its value at its first female-dynast election. Table 19 summarizes the classification.

Table 19: Alignment with the provincial governor, by source of the governor's coalition

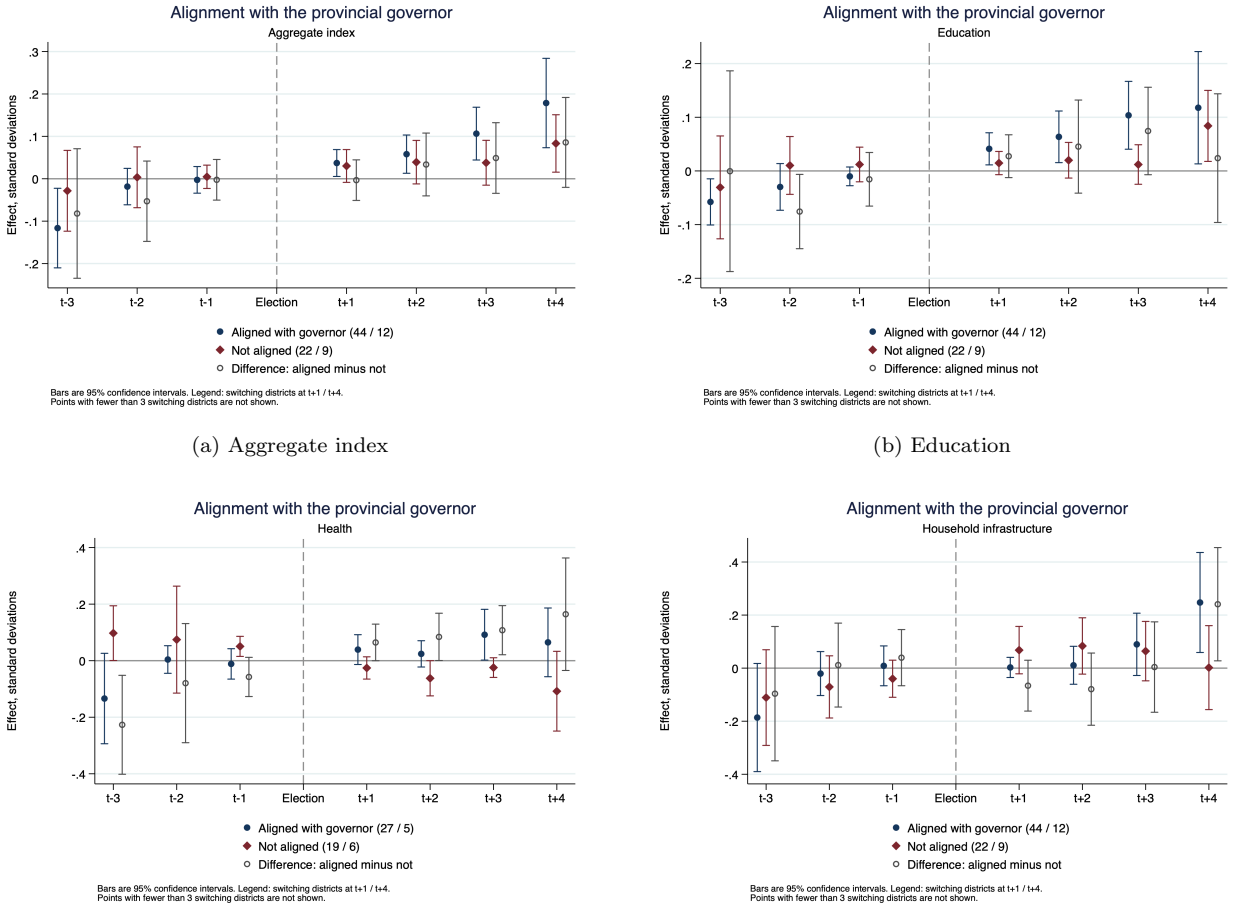
	All district elections		Female-dynast elections	
	Elections	Share aligned	Elections	Share aligned
Directly elected governor	1,221	0.628	90	0.722
Governor elected by the provincial parliament	164	0.488	2	0.500
Not classified	356	–	10	–
Total	1,741	0.612	102	0.717

An election is aligned if the winning ticket shares at least one party with the coalition of the governor in office at the time of the election. The share aligned is computed over classified elections. Elections are not classified when the winning ticket or the governor ran as an independent, when the governor was an acting appointee, or when either coalition cannot be identified. Of the 76 districts that elect a female dynast, 44 are aligned at their first female-dynast election, 23 are not, and 9 are not classified.

Alignment is a party connection, and a district head does not need to come from a dynasty to have it. In the raw data, dynastic tickets are more often in the governor's coalition: 72 percent of female-dynast and male-dynast elections are aligned, against 58 percent of elections won by tickets without a dynast. This gap comes from where and when dynasts win. When I compare elections held in the same province and in the same election round, and therefore under the same governor, female dynasts are no more likely to be aligned than male non-dynasts, with a difference of 0.000 and a standard error of 0.053.

I estimate the effect separately for aligned and unaligned female-dynast districts, with that group as the treatment and every other district as the control, and I run the difference analysis from Section 7 with the alignment indicator. Figure 9 shows the results.

The arm estimates point toward a larger effect in aligned districts. On the aggregate index, the effect in aligned districts rises from 0.037 at the first lag to 0.179 at the fourth, while in unaligned districts it rises from 0.030 to 0.083. The average cumulative effect per treated district, which averages across the lags and is not shown in the figure, is 0.087 (0.025) in aligned districts and 0.050 (0.025) in unaligned districts. The difference analysis puts the gap on the aggregate index between -0.003 and 0.086 across the lags. Education shows



(a) Aggregate index

(b) Education

(c) Health

(d) Household infrastructure

Figure 9: Effect by alignment with the provincial governor

Each colored series is a separate estimation in which the treatment is a district that elects a female dynast whose ticket is (or is not) aligned with the governor's coalition, with every other district as the control. Each point is the cumulative effect at that horizon, and negative horizons are leads. The grey series with hollow markers is the difference analysis: the coefficient on an indicator for alignment in a regression of the district-level effect estimates on that indicator, run over the districts that elect a female dynast. Bars are 95 percent confidence intervals, clustered on district. The legend gives the number of switching districts at the first and fourth lag. Points with fewer than three switching districts are not shown.

the same pattern as the aggregate index. The aligned group rises from 0.041 to 0.118, the unaligned group stays between 0.012 and 0.020 through the third lag, and the difference analysis reaches 0.074 (0.042) at the third lag.

Female dynasts are no more likely than other district heads to be aligned with the governor, so alignment is unlikely to be the channel behind the main results, although the data cannot rule it out.

8.6 Discussion

The mechanism results answer the question that the model leaves open, and they follow its two predictions. The model leaves open whether a family's network helps or hurts the provision of public goods. For female dynasts, the effect is concentrated in families that held office before 2005, so the capacity of an old family outweighs the capture it enables (Section 8.1). What that capacity consists of is the breadth of the network. The effect is largest where her family never held this seat, and these are the women with the most relatives in legislatures and at the provincial level. The model predicts that a family's network helps more when the district head places more weight on delivery. Old families help female dynasts and do not help male dynasts (Section 8.2). The model also predicts that a woman who governs in her own right gains more than a man from the same family, while a woman who governs on her family's behalf gains the same. The gain comes from the women who did not directly succeed a relative, and the women who did look like the men from their families (Section 8.3).

The women who carry the effect bring the same formal experience as the men from the same kind of family, and they did not have to clear a higher bar at the ballot box (Section 8.4). The difference between them and the men from their families therefore does not come from experience on paper or from the voters. Taken together, the results suggest that what a female dynast brings to office is the combination of an old family's network and her own standing. Neither of the two produces the gain on its own: men from old families do not deliver more, and women from newer families look like women without a political family.

Gender norms may affect how quickly a female dynast delivers. In districts with higher female employment, the effect on the aggregate index appears from the first lag, while districts with lower female employment catch up by the later lags (Section 8.5.1). This is consistent with more favorable norms reducing the friction that a female district head meets, although it is also consistent with norms changing who is elected, and the data cannot separate the two. Female employment is also a behavioral proxy that reflects the local labor market as well as norms, and a measure built from direct survey questions on attitudes toward women's work and leadership would allow a sharper test. Alignment with the provincial governor does not appear to be the channel (Section 8.5.2).

There are four limits to these conclusions. The first is statistical precision, and most of the differences between groups in this section are not statistically significant. In the entrenchment split, the difference analysis shows no significant gap at any lag on the aggregate index, education, or household infrastructure, and the health comparison is not informative because the difference analysis has a lead of -0.281 (0.066) at $t-2$. No pairwise difference between

the depth groups is significant. In the comparisons by gender, the difference-in-differences is positive but not significant, at 0.043 (0.041) for entrenchment and 0.054 (0.068) for direct succession, and the difference between the entrenched wives and the other entrenched female dynasts is not significant either. The gaps between women and men within entrenched families and the differences in legislative experience at first win are the exceptions. For gender norms, the difference analysis separates the two groups only at the first lag. For alignment, it shows no significant gap on the aggregate index at any lag and reaches the ten percent level on education only at the third lag. The gaps it finds on health and household infrastructure are the least reliable, because the health comparison has a lead of -0.227 (0.089) at $t-3$, and the household gap appears only at the fourth lag, where it rests on 12 aligned and 9 unaligned districts. What the results show is therefore mostly a consistent ordering of the estimates.

The second limit is the trend before the election. In the entrenchment, gender-norm and alignment splits, the group with the larger effect has a negative lead at $t-3$. On the aggregate index, the lead is -0.102 (0.044) for entrenched families, with a joint placebo p-value of 0.065, and it rests on 14 districts. It is -0.095 (0.040) where female employment is higher, with a joint placebo p-value of 0.001, and -0.216 (0.074) on household infrastructure in the same group. It is -0.116 (0.048) for aligned districts, where it rests on 9 districts. A negative lead means that these districts were already improving relative to the comparison districts before the election, so the estimates for these groups, and the gaps between groups, may be overstated. The main results carry a negative lead at the same horizon (Section 6.5), and these splits may partly be sorting districts by that trend. The shift toward capital spending in Appendix G.2 may also have begun before the election.

The third limit is support. In the entrenchment split, fewer than three districts from newer families are observed at the fourth lag, so the comparison between the two groups rests on the first three lags. The comparisons by gender rely on static two-way fixed effects, because they need several groups in one regression and because the cells become too thin for the heterogeneity-robust estimator to say much at each horizon. The cells with a direct successor hold about 20 districts each, and the first-win comparisons rest on 14 to 19 entrenched women who are not wives for each characteristic.

The fourth limit concerns interpretation. The model's link between direct succession and governing on the family's behalf is an assumption, and the first-win comparisons cannot rule out that the women who carry the effect differ from the men in qualities that a biography does not record.

9 Conclusion

Two literatures address what happens when a district elects a female dynast, and neither gives a clear answer. Female politicians have improved the provision of public goods in many developing countries, and dynastic politicians have more often been found to deliver less, although both sets of findings depend on the setting. I find that electing a female dynast raises the provision of public infrastructure, and that her dynastic ties do not work against this gain. Districts that elect a female dynast gain about 0.07 of a standard deviation of public infrastructure, driven by household infrastructure and education. When I separate the two characteristics, the gender contrast is roughly two and a half times the size of the dynastic one, and female dynasts are the only kind of leader whose effect is distinguishable from zero. Holding gender constant, the effect of the dynastic tie is still positive, and it is smaller than the gender effect without working against it. Neither contrast is estimated precisely enough to rule out that they are equal, so I interpret the comparison only as a direction.

The gains are not uniform. Household infrastructure improves most in less urbanized districts, which had the most room to add basic amenities, and the effect of a deputy arrives later than the effect of a head.

The mechanism results run against the simplest version of the network account. A dynastic family that already controls a district's executive office should, on that account, be the one best placed to deliver, because it knows the bureaucracy and has the relationships to move it. The effect is instead largest in the districts where the family had never held the office before. These women come from families whose networks are wider and older, with more relatives in other offices and a longer presence in the state, and the effect is concentrated where the family first held office before direct elections began. What a female dynast inherits appears to matter more than whether she inherits the seat. Gender norms may also play a role through the selection of the women who win, although a better measure of norms is needed to test this. These results establish an ordering of the estimates. They do not establish significant differences between groups, and the family network can explain at most part of the effect.

What these outcomes can show is also limited by what they measure. The village census records the stock of facilities that exist in a village: how many schools there are, how many clinics there are, and whether most households have safe water. It does not record what people actually receive from those facilities. A district can gain a school building without gaining a teacher to stand in it. I chose the stock because a district head's budget authority operates on the stock, but whether children learned more is a question this paper cannot

answer. Household survey data would answer it, and that is the natural extension of this work.

This paper began with a straightforward question. If a large share of the women who reach office do so as dynasts, does electing them improve outcomes for their constituents? On the outcomes I can measure, it does. Districts that elect a female dynast gain public infrastructure. Most of that gain comes from her being a woman, and the part that comes from her family depends on how wide and how old her family's network is.

References

Afridi, Farzana, Vegard Iversen, and Mamidipudi R Sharan. 2017. "Women Political Leaders, Corruption, and Learning: Evidence from a Large Public Program in India." *Economic Development and Cultural Change*, 66(1): 1–30.

Alesina, Alberto, Paola Giuliano, and Nathan Nunn. 2013. "On the Origins of Gender Roles: Women and the Plough." *Quarterly Journal of Economics*, 128(2): 469–530.

Andersen, Steffen, Erwin Bulte, Uri Gneezy, and John A List. 2008. "Do Women Supply More Public Goods than Men? Preliminary Experimental Evidence from Matrilineal and Patriarchal Societies." *American Economic Review*, 98(2): 376–81.

Asako, Yasushi, Takeshi Iida, Tetsuya Matsubayashi, and Michiko Ueda. 2015. "Dynastic Politicians: Theory and Evidence from Japan." *Japanese Journal of Political Science*, 16(1): 5–32.

Asher, Sam, and Paul Novosad. 2017. "Politics and Local Economic Growth: Evidence from India." *American Economic Journal: Applied Economics*, 9(1): 229–273.

Aspinall, Edward. 2015. "Money Politics: Patronage and Clientelism in Southeast Asia." In *Routledge Handbook of Southeast Asian Democratization.* . 1st ed., , ed. William Case, 299–313. United Kingdom:Routledge, Taylor & Francis Group.

Aspinall, Edward, and Marcus Mietzner. 2010. "Problems of Democratisation in Indonesia: an Overview." In *Problems of Democratisation in Indonesia: Elections, Institutions, and Society.* , ed. Edward Aspinall and Marcus Mietzner, 1–17. Singapore:Institute of Southeast Asian Studies. Section: 1.

Aspinall, Edward, and Wawan Mas'udi. 2017. "The 2017 Pilkada (Local Elections) in Indonesia: Clientelism, Programmatic Politics and Social Networks." *Contemporary Southeast Asia: A Journal of International and Strategic Affairs*, 39(3): 417–426.

Badan Pusat Statistik. 2021*a*. "Hasil Sensus Penduduk 2020." Badan Pusat Statistik Berita Resmi Statistik, Jakarta.

Badan Pusat Statistik. 2021*b*. "Peraturan Kepala Badan Pusat Statistik Nomor 120 Tahun 2020 tentang Klasifikasi Desa Perkotaan dan Perdesaan di Indonesia 2020." Badan Pusat Statistik, Jakarta.

Badan Pusat Statistik. 2021*c*. "Pertumbuhan Ekonomi Indonesia Triwulan IV-2020." Badan Pusat Statistik Berita Resmi Statistik, Jakarta.

Bagues, Manuel, and Pamela Campa. 2021. "Can Gender Quotas in Candidate Lists Empower Women? Evidence from a Regression Discontinuity Design." *Journal of Public Economics*, 194: 104315.

Baskaran, Thushyanthan, and Zohal Hessami. 2019. "Competitively Elected Women as Policy Makers." CESifo Working Paper.

Beaman, Lori, Raghabendra Chattopadhyay, Esther Duflo, Rohini Pande, and Petia Topalova. 2009. "Powerful Women: Does Exposure Reduce Bias?" *The Quarterly Journal of Economics*, 124(4): 1497–1540.

Besley, Timothy, and Stephen Coate. 1997. "An Economic Model of Representative Democracy." *The Quarterly Journal of Economics*, 112(1): 85–114.

Besley, Timothy, Olle Folke, Torsten Persson, and Johanna Rickne. 2017. "Gender Quotas and the Crisis of the Mediocre Man: Theory and Evidence from Sweden." *American Economic Review*, 107(8): 2204–2242.

Bhalotra, Sonia, and Irma Clots-Figueras. 2014. "Health and the Political Agency of Women." *American Economic Journal: Economic Policy*, 6(2): 164–197.

Bhalotra, Sonia, Damian Clarke, Joseph Flavian Gomes, and Atheendar Venkataramani. 2023. "Maternal Mortality and Women's Political Power." *Journal of the European Economic Association*, jvad012.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2023. "Revisiting Event Study Designs: Robust and Efficient Estimation." *Review of Economic Studies, forthcoming.*

Bragança, Arthur, Claudio Ferraz, and Juan Rios. 2015. "Political Dynasties and the Quality of Government." *Unpublished manuscript.*

Brollo, Fernanda, and Tommaso Nannicini. 2012. "Tying Your Enemy's Hands in Close Races: The Politics of Federal Transfers in Brazil." *American Political Science Review*, 106(4): 742–761.

Brollo, Fernanda, and Ugo Troiano. 2016. "What Happens When a Woman Wins an Election? Evidence from Close Races in Brazil." *Journal of Development Economics*, 122: 28–45.

Callaway, Brantly, and Pedro HC Sant'Anna. 2021. "Difference-in-differences with Multiple Time Periods." *Journal of Econometrics*, 225(2): 200–230.

Chattopadhyay, Raghabendra, and Esther Duflo. 2004. "Women as Policy Makers: Evidence From a Randomized Policy Experiment in India." *Econometrica*, 72(5): 1409–1443.

Cirone, Alexandra, and Thomas B. Pepinsky. 2022. "Historical Persistence." *Annual Review of Political Science*, 25: 241–259.

CNN Indonesia. 2017. "Bupati Sri Hartini Divonis 11 Tahun Penjara." 20 September 2017.

Cruz, Cesi, Julien Labonne, and Pablo Querubin. 2017. "Politician Family Networks and Electoral Outcomes: Evidence from the Philippines." *American Economic Review*, 107(10): 3006–37.

Dal Bó, Ernesto, Pedro Dal Bó, and Jason Snyder. 2009. "Political Dynasties." *The Review of Economic Studies*, 76(1): 115–142.

Dar, Aaditya. 2019. "Parachuters vs. Climbers: Economic Consequences of Barriers to Political Entry in a Democracy." *Job Market Paper*.

De Chaisemartin, Clément, and Xavier d'Haultfoeuille. 2020. "Two-way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review*, 110(9): 2964–2996.

Deichmann, Uwe, Kai Kaiser, Somik V. Lall, and Zmarak Shalizi. 2005. "Agglomeration, Transport, and Regional Development in Indonesia." The World Bank Policy Research Working Paper 3477.

Dell, Melissa, and Benjamin A. Olken. 2020. "The Development Effects of the Extractive Colonial Economy: The Dutch Cultivation System in Java." *Review of Economic Studies*, 87(1): 164–203.

Dollar, David, Raymond Fisman, and Roberta Gatti. 2001. "Are Women Really the "Fairer" Sex? Corruption and Women in Government." *Journal of Economic Behavior & Organization*, 46(4): 423–429.

Downs, Anthony. 1957. "An Economic Theory of Democracy." *Harper and Row*, 28.

Ferreira, Fernando, and Joseph Gyourko. 2014. "Does Gender Matter for Political Leadership? The Case of U.S. Mayors." *Journal of Public Economics*, 112: 24–39.

Folke, Olle, Johanna Rickne, and Daniel M Smith. 2021. "Gender and Dynastic Political Selection." *Comparative Political Studies*, 54(2): 339–371.

Gagliarducci, Stefano, and M. Daniele Paserman. 2012. "Gender Interactions within Hierarchies: Evidence from the Political Arena." *Review of Economic Studies*, 79(3): 1021–1052.

George, Siddharth Eapen. 2020. "Like Father, Like Son? The Effect of Political Dynasties on Economic Development." *Work. Pap., Harvard Univ., Cambridge, MA Google Scholar Article Location*, 4.

Geys, Benny, and Daniel M Smith. 2017. "Political Dynasties in Democracies: Causes, Consequences and Remaining Puzzles." *The Economic Journal*, 127: F446–F454.

Government of Indonesia. 1974. "Undang-Undang No. 5 tahun 1974 Tentang Pokok-Pokok Pemerintahan di Daerah." (Law No. 5 of 1974 Concerning the Principles of Regional Government), art. 16.

Government of Indonesia. 1999*a*. "Undang-Undang No. 22 tahun 1999 Tentang Pemerintahan Daerah." (Law No. 22 of 1999 Concerning Regional Government), arts. 18, 34–40.

Government of Indonesia. 1999*b*. "Undang-Undang No. 3 tahun 1999 Tentang Pemilihan Umum." (Law No. 3 of 1999 Concerning General Elections), arts. 1(7), 41(5), 68.

Government of Indonesia. 2003. "Undang-Undang No. 12 tahun 2003 Tentang Pemilihan Umum Anggota Dewan Perwakilan Rakyat, Dewan Perwakilan Daerah, dan Dewan Perwakilan Rakyat Daerah." (Law No. 12 of 2003 Concerning the Election of Members of the DPR, DPD and DPRD), art. 107(2).

Government of Indonesia. 2014. "Undang-Undang No. 23 tahun 2014 Tentang Pemerintahan Daerah." Published: (Law No. 23 of 2014 Concerning Local Government).

Hessami, Zohal, and Mariana Lopes da Fonseca. 2020. "Female Political Representation and Substantive Effects on Policies: A Literature Review." *European Journal of Political Economy*, 63: 101896.

Hsiao, Allan. 2019. "Democratization and Infrastructure Investment: Evidence from Healthcare in Indonesia." *Working paper*.

Jakiela, Pamela. 2021. "Simple Diagnostics for Two-way Fixed Effects." *arXiv preprint arXiv:2103.13229*.

Kenawas, Yoes C. 2020. "Dynastic Politics in Indonesia."

Kenawas, Yoes. C. 2022. "Indonesian Political Dynasty Dataset."

Kis-Katos, Krisztina, and Bambang Suharnoko Sjahrir. 2017. "The Impact of Fiscal and Political Decentralization on Local Public Investment in Indonesia." *Journal of Comparative Economics*, 45(2): 344–365.

Labonne, Julien, Sahar Parsa, and Pablo Querubin. 2021. "Political Dynasties, Term Limits and Female Political Representation: Evidence from the Philippines." *Journal of Economic Behavior & Organization*, 182: 212–228.

Lippmann, Quentin. 2022. "Gender and Lawmaking in Times of Quotas." *Journal of Public Economics*, 207: 104610.

Mahkamah Konstitusi Republik Indonesia. 2008. "Putusan Nomor 22-24/PUU-VI/2008." (Constitutional Court Decision No. 22-24/PUU-VI/2008, 23 December 2008), on Article 214 of Law No. 10 of 2008.

Malik, Adeel, Rinchan Ali Mirza, and Jean-Philippe Platteau. 2021. "Entrenched Political Dynasties and Development under Competitive Clientelism: Evidence from Pakistan." Economic Development and Institutions EDI Working Paper.

Martinez-Bravo, Monica, Priya Mukherjee, and Andreas Stegmann. 2017. "The Non-democratic Roots of Elite Capture: Evidence from Soeharto Mayors in Indonesia." *Econometrica*, 85(6): 1991–2010.

Mietzner, Marcus. 2016. "The Sukarno Dynasty in Indonesia: Between Institutionalisation, Ideological Continuity and Crises of Succession." *South East Asia Research*, 24(3): 355–368.

Mietzner, Marcus. 2020. "Indonesian Parties Revisited: Systemic Exclusivism, Electoral Personalisation and Declining Intraparty Democracy." In *Democracy in Indonesia: From stagnation to regression?*. , ed. Thomas Power and Eve Warburton, 191–209. Singapore:ISEAS - Yusof Ishak Institute. Section: 10.

Osborne, Martin J, and Al Slivinski. 1996. "A Model of Political Competition with Citizen-Candidates." *The Quarterly Journal of Economics*, 111(1): 65–96.

Pande, Rohini. 2003. "Can Mandated Political Representation Increase Policy Influence for Disadvantaged Minorities? Theory and Evidence from India." *American economic review*, 93(4): 1132–1151.

Prihatini, Ella, and Iim Halimatusa'diyah. 2022. "Gender, Political Dynasties, and Committee Assignments: Evidence From Indonesia." *Parliamentary Affairs*, gsac019.

Querubin, Pablo. 2016. "Family and Politics: Dynastic Persistence in the Philippines." *Quarterly Journal of Political Science*, 11(2): 151–181.

Ricklefs, M. C. 2008. *A History of Modern Indonesia since c. 1200.* . 4 ed., Stanford, CA:Stanford University Press.

Sherlock, Stephen. 2004. "The 2004 Indonesian Elections: How the System Works and What the Parties Stand For." *Canberra: Centre for Democratic Institutions, Research School of Social Sciences, Australian National University.*

Smith, Daniel M, and Shane Martin. 2017. "Political Dynasties and the Selection of Cabinet Ministers." *Legislative Studies Quarterly*, 42(1): 131–165.

Smith, Daniel Markham. 2012. "Succeeding in Politics: Dynasties in Democracies." *Doctoral Dissertation.*

Sun, Liyang, and Sarah Abraham. 2021. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." *Journal of Econometrics*, 225(2): 175–199.

Tirto.id. 2016. "Sri Hartini dan Dinasti Kekuasaan di Klaten." 30 December 2016.

Zhu, Yifei, Songpo Yang, Jiangnan Zhu, and Junyan Jiang. 2026. "Agentic Framework for Political Biography Extraction." Working paper.

Appendix

A. A simple model of a district head's payoff

This appendix sets out the model behind the conceptual framework in Section 2.

A.1. Setup

A district head controls a budget $B > 0$ and chooses a share $s \in [0, 1]$ of it to divert to private rents for herself or her family. The rest is spent on public goods. Her family's network is $N \geq 0$, which I think of as how long and how deeply the family has been entrenched in office. The network does two things. It raises the amount of public goods that each unit of the budget delivers, $a(N) > 0$ with $a'(N) > 0$, because the family has ties to the bureaucracy, the local parliament, the parties, and higher levels of government. It also raises the private return to diverting the budget, $v(N) > 0$ with $v'(N) > 0$, because an entrenched family has more power to extract rents and faces less electoral threat. Delivered public goods are

$$G = a(N)\,(1 - s)\,B. \tag{11}$$

The district head places a weight $w > 0$ on delivered public goods. This weight captures her own preferences over public goods and the value she places on building a reputation of her own. Diverting the budget is costly, for example because of the risk of detection or of losing office, and the cost is convex in the share diverted. Her payoff is

$$U(s) = w\,a(N)\,(1 - s)\,B \;+\; v(N)\,s\,B \;-\; \frac{k}{2}\,s^2 B, \qquad k > 0. \tag{12}$$

A.2. Solution

The payoff is strictly concave in s. The first-order condition is

$$-\,w\,a(N)\,B + v(N)\,B - k\,s\,B = 0, \tag{13}$$

so the optimal share diverted is

$$s^* = \frac{v(N) - w\,a(N)}{k}. \tag{14}$$

I assume $0 < v(N) - w\,a(N) < k$, so that the solution is interior: the district head diverts some of the budget, but not all of it. The share diverted rises with the private return to

rents and falls with the weight she places on delivery and with the capacity of the network. Substituting s^* into G gives delivered public goods,

$$G^* = a(N)\,B\left[1 - \frac{v(N) - w\,a(N)}{k}\right] = a(N)\,B - \frac{a(N)\,v(N)\,B}{k} + \frac{w\,a(N)^2\,B}{k}. \tag{15}$$

A.3. Results

Result 1: the effect of the family network is ambiguous. Differentiating Equation (15) with respect to N,

$$\frac{\partial G^*}{\partial N} = \frac{B}{k}\Big[\,a'(N)\big(k - v(N)\big) + 2\,w\,a(N)\,a'(N) \;-\; a(N)\,v'(N)\,\Big]. \tag{16}$$

The first two terms are the capacity force: a stronger network delivers more public goods for each unit of the budget, both directly and because a district head who values delivery diverts less when each unit delivers more. Together they equal $a'(N)\big[k - \big(v(N) - w\,a(N)\big) + w\,a(N)\big]$, which is positive under the interior condition. The last term is negative. It is the capture force: a stronger network makes diversion more attractive. The sign of $\partial G^*/\partial N$ depends on which of the two forces is larger, so the effect of the family network is an empirical question.

Result 2: the family network helps more when the district head places more weight on delivery. The cross-partial derivative is

$$\frac{\partial^2 G^*}{\partial N\,\partial w} = \frac{2\,a(N)\,a'(N)\,B}{k} > 0. \tag{17}$$

A district head who places more weight on delivery diverts less, so more of the capacity that her family brings reaches her constituents. The capture force does not depend on w, so a higher weight shifts the balance in Equation (16) toward capacity.

Result 3: own standing and proxy. Let w_M be the weight that a male district head places on delivery, and let $w_F > w_M$ be the weight of a woman who governs in her own right. A woman who governs on behalf of her husband or family takes on the family's weight, w_M. The difference between a woman who governs in her own right and a man from a family with the same network is

$$G^*(w_F, N) - G^*(w_M, N) = \frac{(w_F - w_M)\,a(N)^2\,B}{k} > 0, \tag{18}$$

which grows with N because $a'(N) > 0$. A woman who governs as a proxy delivers $G^*(w_M, N)$, the same as a man from her family.

A.4. Mapping to the empirical comparisons

The comparison between entrenched and newer families tells us the sign of Result 1 for female dynasts. The comparison between female and male dynasts from entrenched families tests Result 2. The comparison between women who directly succeeded a relative and women who did not tests Result 3, under the assumption that direct succession is the case in which governing on the family's behalf is most plausible. The model does not derive that assumption.

A.5. What the model leaves out

There are three things that the model leaves out. First, it takes the district head as given and says nothing about who is selected into office, so the evidence on characteristics at first election in Section 8 complements the model and does not follow from it. Second, the weight w is reduced form. It stands for two foundations, her own preferences over public goods and the value of building a reputation of her own, and the model does not separate them. Third, the model does not include gender norms. One way to add them is as a friction that lowers a woman's capacity, $a_F(N) = a(N) - \phi$, where better gender norms lower ϕ.

B. Election and dynasty data

This appendix describes how I built the dataset of local elections and dynastic status that I introduce in Section 4.1. The dataset covers the direct local elections from 2005 to 2024 for governors, *bupati* and *walikota*, together with their deputies. It contains 2,394 elections.

B.1. Sources for each election cycle

No single source covers every election cycle, so the base data come from a different place in each period. Table A1 summarizes them.

For 2005 to 2008, I start from the election data that Martinez-Bravo, Mukherjee and Stegmann (2017) assembled, which record the winning pairs, the vote counts and the party coalitions. For 2010, I start from the candidate and vote tables that the General Elections Commission (KPU) published. In both periods I completed the missing elections, candidates, votes and coalitions by hand. For 2011, 2012 and 2013 there is no consolidated KPU record, so I collected these elections myself from news reporting, from Constitutional Court decisions on election disputes, and from local government records. For 2015, 2017 and 2018, I build on the data that Kenawas (2022) collected. His data record the winning pairs, the other

candidate pairs and the dynastic status of the winning head and deputy, and he shared them with me in 2023. The election and vote data for these three rounds are my own work. For 2020 and 2024, the KPU published tables that list every candidate pair with its votes, its party coalition and the gender of both candidates. I extracted the elections for those two rounds from these tables.

Table A1: Sources of the election and dynasty data, by election cycle

Cycle	Elections, candidates and votes	Dynastic status, winning head and deputy	Dynastic status, other candidates
2005–2008	Martinez-Bravo, Mukherjee and Stegmann (2017), completed by hand	Hand-coded	AI agent
2010	KPU, completed by hand	Hand-coded	AI agent
2011–2013	Collected by hand	Hand-coded	AI agent
2015, 2017, 2018	Kenawas (2022); vote data collected by hand	Kenawas (2022), rechecked by AI agent	Hand-coded
2020, 2024	KPU, extracted by AI agent	AI agent	AI agent

In every cycle, the AI agent also filled in elections and candidate pairs that were missing from the base data. Section B.3 describes the protocol that the agent followed.

B.2. What counts as a dynastic tie

Section 4.1 gives the definition of a dynastic politician. This subsection describes how I apply it.

A relative's office qualifies if it is one of the following: president or vice president, governor, *bupati* or *walikota*, or a seat in the national parliament (DPR), the Regional Representative Council (DPD), the People's Consultative Assembly (MPR), or a provincial or district parliament (DPRD). Before direct elections began in 2005, governors and district heads were chosen by the regional parliaments or appointed by the central government. I count these heads as officeholders, because each of them was the definitive holder of the office for a term.

Several positions resemble political office and do not qualify. I do not count acting heads (*penjabat* and *pelaksana tugas*), who are caretakers appointed to fill the office between elections. I also do not count the governor of Yogyakarta, who holds the office under the special-region law and is not elected. Cabinet ministers, party officials, military and police officers, civil servants and village heads do not qualify either. Neither do religious leaders, holders of traditional or royal titles, and the leaders of mass organizations. A relative who ran for office and lost does not qualify.

I count relatives by blood out to grandparents, grandchildren, aunts, uncles and cousins. I count relatives by marriage when they are a spouse, a parent-in-law, a child-in-law or a sibling-in-law. A former spouse or a late spouse counts if that person held a qualifying office. A shared surname or clan name is not evidence of a tie, and a political partnership such as a shared ticket is not evidence of a tie either.

I code dynastic status as of each election. A candidate is a descendant if a relative held a qualifying office before her. A candidate is a founder if she held office first and a relative entered office after her. Because a founder is identified by what her relatives do later, I searched for relatives' offices up to and including the 2024 legislative and local elections. A losing candidate who has never held office does not become a dynast at an election because a relative enters office after it.

For each relative, I record the office in which that relative first served. This identifies which member of the family entered politics first. It also means that the data understate the offices that relatives went on to hold later.

B.3. How I coded dynastic status

I hand-coded the dynastic status of the winning heads and deputies in every election from 2005 to 2013. I also hand-coded the dynastic status of the losing candidates in 2015, 2017 and 2018. For the winners of those three rounds, I use the codes of Kenawas (2022). My sources were Indonesian Wikipedia, national and regional news outlets, candidate profiles published by the KPU and by the parliaments, and Constitutional Court decisions.

I then used an AI agent to complete the database.[4] The agent did four things. First, it filled in elections and candidate pairs that were missing from 2005 to 2018. Second, it coded the dynastic status of the losing candidates from 2005 to 2013. Third, it extracted the 2020 and 2024 elections from the KPU tables and coded the dynastic status of every candidate in those two rounds. Fourth, it searched again for relatives of the 2015, 2017 and 2018 candidates who entered office after Kenawas (2022) collected his data. This search changed 57 winning heads and deputies from non-dynasts to founders.

The agent followed the synthesis-then-coding framework of Zhu et al. (2026). In this framework, a supervisor gives the agent a self-contained task. The agent searches the web, filters what it finds, and consolidates the evidence, and a coding step then turns that evidence into structured records. I acted as the supervisor. I wrote the task instructions, which state the research goal, the coding rules above, the sources to search first, and a limit on the number of searches per person. The agent searched up to three times for each person, and

[4]The agent was Claude Code, which runs on Anthropic's Claude models.

it coded a person as a non-dynast if it found no relative in a qualifying office after three searches. Before the agent recorded a candidate as a dynast, it confirmed the tie in a second source, usually Wikipedia. The agent flagged every case that it could not resolve, for example a shared surname with no documented tie or a relative whose office could not be dated. I adjudicated each of these cases by hand.

Of the 12,266 dynastic-status codes in the data, about 31 percent are my own hand-coding. About 9 percent come from Kenawas (2022), and about 60 percent come from the agent's completion of the database. For the winning heads and deputies from 2005 to 2018, the shares are 52 percent, 34 percent and 14 percent. For 2020 and 2024, the agent coded every candidate.

B.4. Checks

I checked the completed data in three ways. First, I compared the election results against Constitutional Court decisions and other sources in 43 rounds of audit, and I resolved each disagreement by hand. Second, I had the agent search again for every 2024 winner coded as a non-dynast, with the adult children of the candidate as the main target of the search. Third, I compared the agent's codes for 2015, 2017 and 2018 against the codes of Kenawas (2022). The only changes were the 57 new founders, and no dynast was changed to a non-dynast.

B.5. Limitations

There are three limitations of the data. The first is that the early elections are incomplete. For some elections between 2005 and 2013, mostly in Papua, in Maluku and in newly created districts, I could recover the winning pair but not the other candidates or the vote counts. The second is that a code of zero means that no qualifying relative was found after three searches. For losing candidates with little public record, this is weaker evidence than it is for winners, whose careers are better documented. The third is that founder status depends on what relatives do after the election. A founder whose relative enters office after 2024 appears in the data as a non-dynast.

B.6. Biographies of winning heads and deputies

I collected biographical information on the winning heads and deputies of district elections from 2005 to 2020 who belong to three groups: female dynastic descendants, male dynastic descendants, and female non-dynasts. I did not collect it for male non-dynasts, who are the

large majority of winners, or for founders. The data contain 436 person-elections and 342 individuals.

There are four fields. The first two, year of birth and education, are recorded once for each person. I record education as the highest level completed before the person first won a district election, so a degree counts only if it was completed before that first win. A degree with no date is left as a lower bound or as missing. The other two fields are recorded for each election. The position held before the election is the most recent role the person held before the campaign. When the person held two roles at once, I record the more formal office. Prior public office records whether the person held any public office before the election, and which. Public office includes a seat in a legislature at any level, the post of head or deputy head of a region at any level, a ministerial or other political appointment, an acting headship, and the post of village head. A career civil servant does not hold public office in this sense. For each office I also record how the person obtained it: by popular vote, by a vote of the regional parliament, as a replacement chosen by the regional parliament, by succession, by appointment, or by a village election. A person won a popular vote before if any of these offices was obtained by popular vote.

The same AI agent that completed the dynasty data collected these fields, under a written protocol that I supervised. The sources are Indonesian and English Wikipedia, the candidate records that the General Elections Commission (KPU) published for the 2024 regional elections, the KPU records of candidates in legislative elections, news reports, regional government websites, and archived candidate records from the 2015 regional elections. The agent coded a value only from the text of a source page that it had read, and never from a search engine's summary. It stopped after three searches for a field, and a field with no source stays missing. Pages from personal blogs or campaign sites were accepted only when the page itself stated the fact. Every match between a source and a person was checked on name, gender, and district or date of birth, and namesakes were rejected. The KPU legislative candidate records were used for year of birth only, because they often list only the minimum level of education required to run.

I reviewed by hand the 20 cases that the agent could not resolve, including conflicting sources and cases where the existing data appeared to be wrong. The existing data also recorded a family relationship in place of a career for some winners, for example "wife of the head", and I replaced these with the role that the person held.

Table A2 reports coverage. Year of birth is known for 386 of the 436 person-elections, the position held before the election for 367, prior public office for 393, and education at first win for 247. Education is the weakest field, mostly because many sources give a degree without

the year in which it was completed. A value of zero for prior public office that rests on a short Wikipedia page is weaker evidence than a zero from a full biography, and I flag these cases in the data.

Table A2: Coverage of the biography data, person-elections

Field	Known
Year of birth	386
Position held before the election	367
Prior public office	393
Education at first win	247
Person-elections	436
Individuals	342

Winning heads and deputies of district elections, 2005–2020, who are female dynastic descendants, male dynastic descendants, or female non-dynasts. Counts are from `winner_biographies_v1_261007.dta`.

C. Robustness of the main effects

Table A3 reports the checks described in Section 6.4 across all four outcomes. Every estimate is the robust estimator's average cumulative effect per treated district, so the rows are directly comparable with the second pair of columns in Table 7.

Table A3: Robustness of the main effects

Specification	Aggregate	Education	Health	Household
Main specification	0.0747***	0.0574***	0.0222	0.0728**
	(0.0175)	(0.0190)	(0.0191)	(0.0334)
Unweighted	0.0721**	0.0878**	0.0306	0.0454
	(0.0338)	(0.0428)	(0.0738)	(0.0355)
Including ex-ante founders	0.0726***	0.0623***	0.0244	0.0576*
	(0.0170)	(0.0186)	(0.0182)	(0.0329)

Each entry is the average cumulative effect per treated district for the outcome in that column. Estimates come from the heterogeneity-robust estimator of De Chaisemartin and d'Haultfoeuille (2020), on the district-by-wave panel, with districts weighted by their village count. Standard errors, in parentheses beneath each estimate, are clustered on district. Significance: *** 1%, ** 5%, * 10%. The main specification weights each district by its village count and defines treatment as the election of a female dynast as head or deputy. The unweighted row is the same specification with the village weights removed. The last row widens treatment to include women who already had a relative in office at the time of their own election.

One of the checks in Section 6.4 is not in this table. The sub-district panel is a separate estimation on a different unit and a shorter window, so its estimates do not belong in a table

whose rows are specifications of the same regression, and I report them in the text instead. I also estimated a village-level panel. It begins in 2011 and drops the three earliest waves, which makes it a different sample and not a different specification, so I do not report it.

D. Difference regressions, all splits and all outcomes

Appendix Figures A1 to A3 report the difference regression for each of the three splits across all four outcomes. Nothing in the aggregate, education or health panels is significant in the effects for any split. The only significant coefficients outside the household panel are placebos.

Figure A1: Urbanization as a predictor, all outcomes

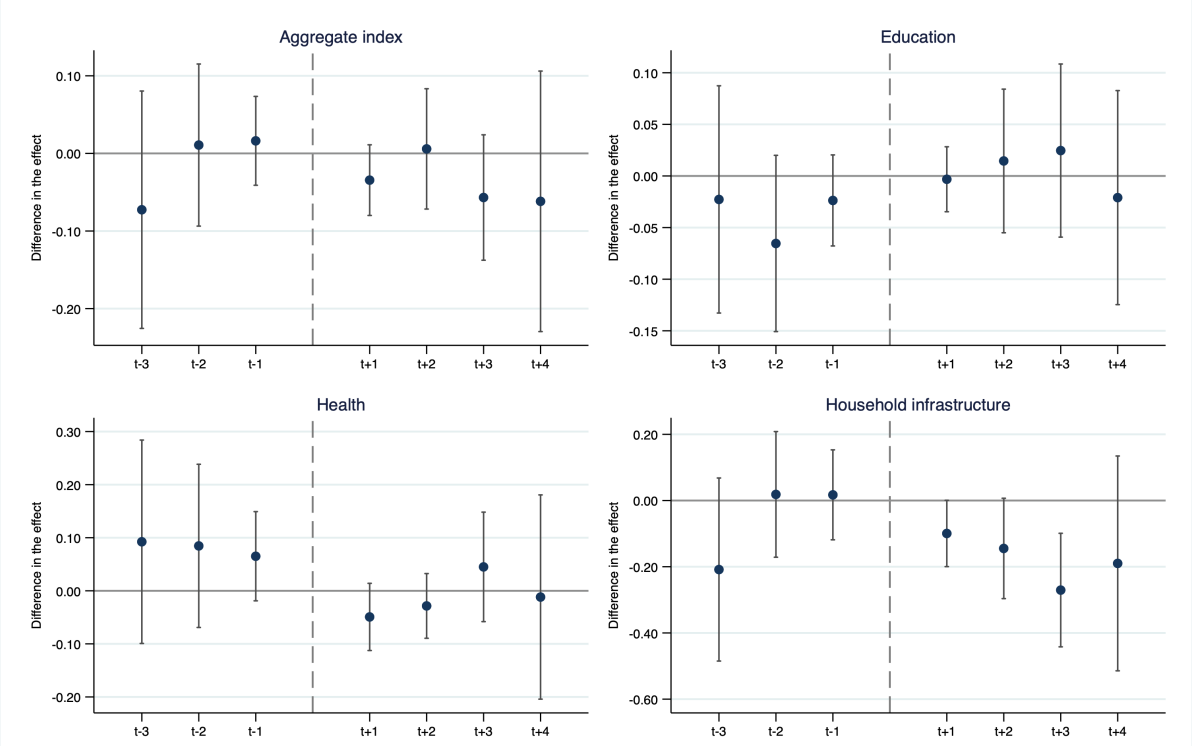
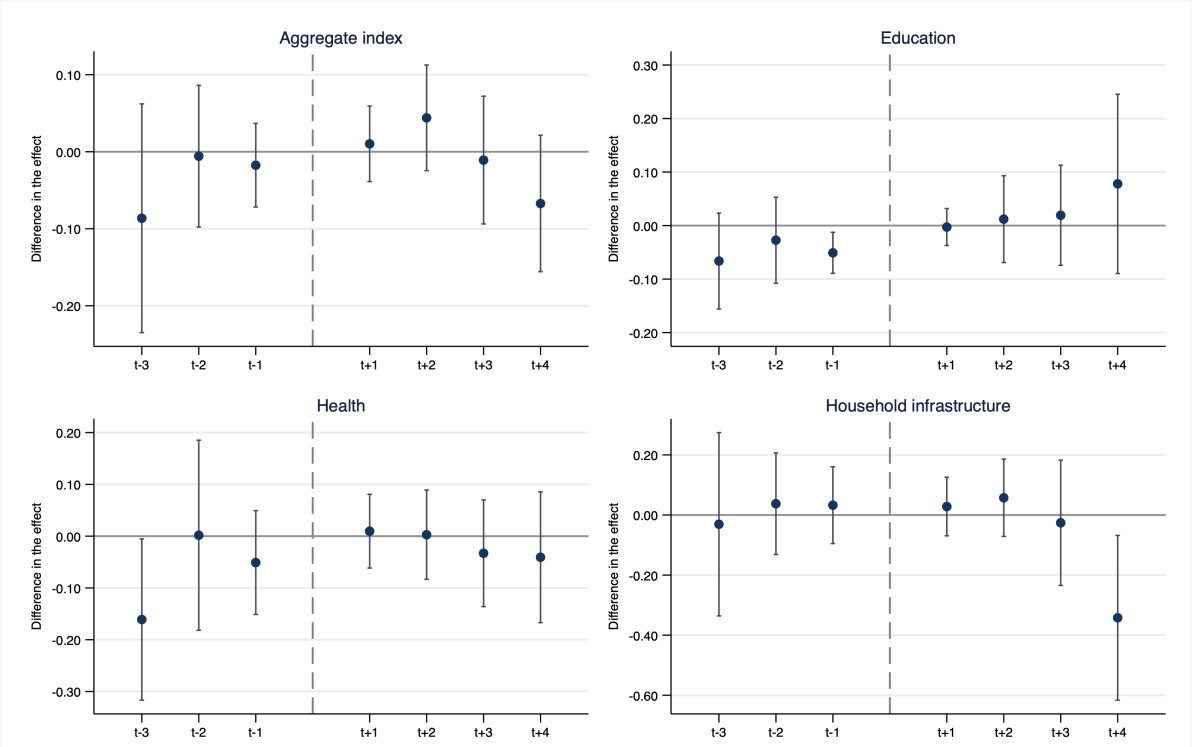


Figure A2: Position on the ticket as a predictor, all outcomes



E. Supplementary tables for the depth and breadth of the network

Table A4 reports how the 102 treated elections divide across the three groups in Section 8.1.1. Table A5 reports the no-prior-claim group's effect on household infrastructure, and Table A6 reports the no-prior-claim group's budget results in levels.

Table A4: The three predecessor classification groups

Group	Definition	Elections	Districts
Direct handover	Family held the office directly before her	41	41
Delayed handover	Family held the office, a term intervened	31	24
No prior claim	Family never held this office	29	23
Total classified		101	

Classification of the 102 treated elections by whether the politician's family had previously held the district executive office she won. Entries are counts, not estimates.

Figure A3: Java as a predictor, all outcomes

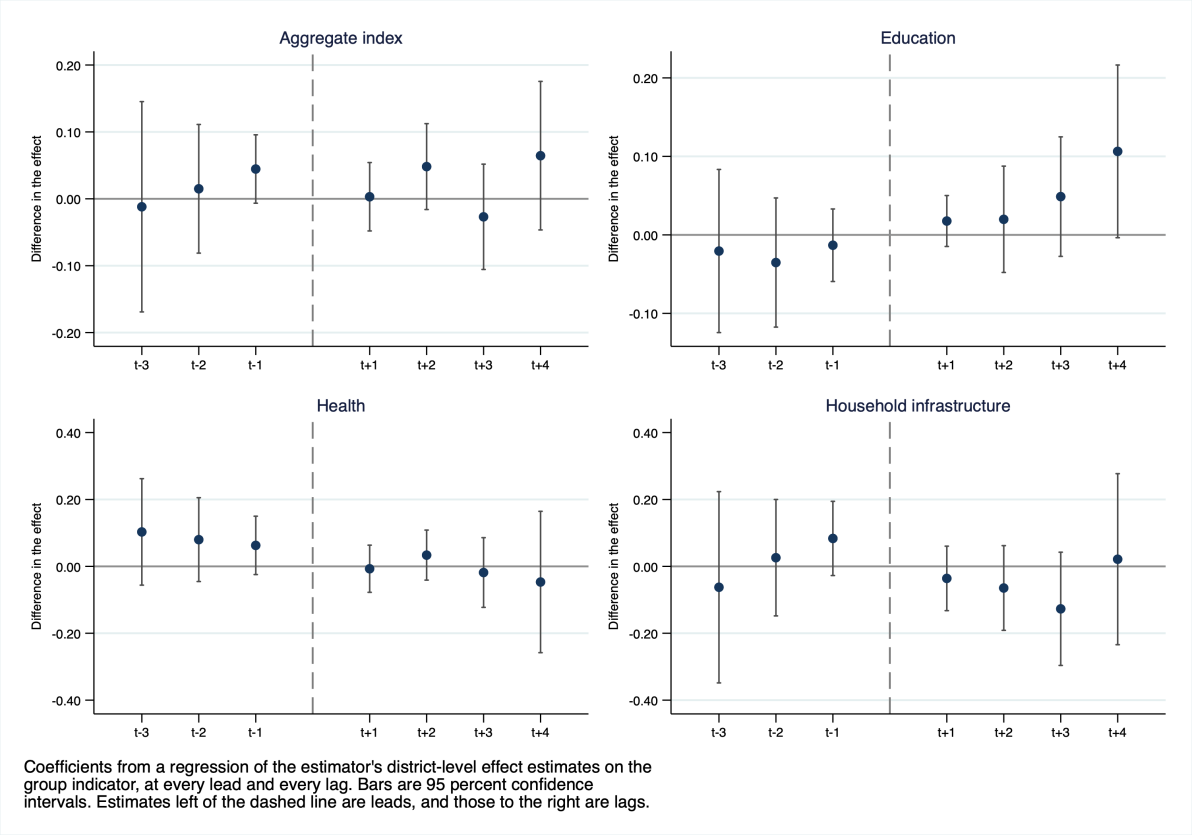


Table A5: No prior claim, household infrastructure

Specification	Estimate	Joint placebo p-value	Switchers h_1 / h_4
Main	0.1416***	0.310	22 / 10
	(0.0438)		
Disagreeing districts dropped	0.1541***	0.251	18 / 7
	(0.0317)		

Entries are the average cumulative effect per treated district on household infrastructure for the no-prior-claim group. The second row drops the districts whose treated elections do not agree with each other on the classification. Estimates come from the heterogeneity-robust estimator of De Chaisemartin and d'Haultfoeuille (2020), on the district-by-wave panel, with districts weighted by their village count. Standard errors, in parentheses beneath each estimate, are clustered on district. Significance: *** 1%, ** 5%, * 10%. The joint placebo column is a nullity test on the lead coefficients. Switcher counts are at the first and fourth lags.

Table A6: No prior claim, budget in levels

	Capital	Personnel	Goods & services	Other	Total expenditure
Estimate	0.1743**	-0.0036	-0.0425	0.2333*	0.0524**
	(0.0746)	(0.0326)	(0.0558)	(0.1283)	(0.0266)
Joint placebo p	0.1396	0.1575	0.1264	0.0105	0.0002

Entries are the average cumulative effect for the no-prior-claim group on the log of each spending category, estimated on the annual district budget panel, 2003–2020. Standard errors, in parentheses beneath each estimate, are clustered on district. Significance: *** 1%, ** 5%, * 10%. The joint placebo row is a nullity test on the lead coefficients. The other category and total expenditure both fail it.

F. Robustness of the ever-officeholder and no-officeholder gap

Table A7 reports the binary split of Section 8.1.1 under four alternative specifications, alongside the main one. The first drops the districts whose treated elections do not agree with each other on the classification. The second removes the other group from the control pool. The third drops the village-count weights. The fourth drops the districts whose only treated election happened before 2015. The gap between the two groups survives all four. It widens in three of them and narrows slightly in one, and there is no specification in which the no-officeholder group falls below the ever-officeholder group.

There are two rows in this table that should be read with care. The unweighted specification produces the widest gap between the two groups, but it also produces the least precise estimates of any specification I run, so I do not read it as evidence that the true gap is larger than the main specification says it is. Dropping the districts that were treated only before 2015 removes treated units from the sample and changes the timing structure of the panel, so this specification estimates a different quantity from the main one. I report it because it shows that the gap is not an artefact of the earliest cohorts. It is not a specification I prefer to the main one.

G. Depth of the family network and the district budget

This appendix reports the estimates behind the depth comparison in Section 8.1.1 and a check on how the breadth of the family network works through the district budget.

Table A7: Robustness of the ever-officeholder / no-officeholder gap

	Family ever held the office		Family never held it	
Specification	Estimate	h_1/h_4	Estimate	h_1/h_4
Main	0.0551***	56/15	0.1175***	22/10
	(0.0184)		(0.0376)	
Districts with disagreeing elections dropped	0.0566***	52/14	0.1214***	18/ 7
	(0.0197)		(0.0411)	
Restricted control pool	0.0601***	52/14	0.1273***	18/ 7
	(0.0198)		(0.0412)	
Unweighted	0.0332	56/15	0.1893**	22/10
	(0.0255)		(0.0954)	
Districts treated only before 2015 dropped	0.0580***	50/ 9	0.1139**	19/ 7
	(0.0151)		(0.0479)	

Entries are the average cumulative effect per treated district on the aggregate index, under the binary split on whether the family ever held the office, so the direct-handover and delayed-handover groups are pooled into the ever-officeholder group. Estimates come from the heterogeneity-robust estimator of De Chaisemartin and d'Haultfoeuille (2020), on the district-by-wave panel, with districts weighted by their village count. Standard errors, in parentheses beneath each estimate, are clustered on district. Significance: *** 1%, ** 5%, * 10%. Switcher counts are the number of treated districts contributing at horizons 1 and 4. The unweighted row drops the village-count weights. The last row removes treated units and alters the timing structure of the panel, so it estimates a different quantity from the main specification.

G.1. Depth of the family network

To measure depth, I match the relatives recorded in the dynasty data against each district's roster of past heads and deputies (Section 4.2), and I sort the 102 treated elections into three groups. In the direct-handover group, her family held the office immediately before her (41 elections in 41 districts). In the delayed-handover group, her family held the office in the past and at least one term passed between that relative and her (31 elections in 24 districts). In the no-prior-claim group, she is the first in her family to hold the executive office of that district (29 elections in 23 districts). I estimate the effect separately for each group, with that group as the treatment and every other district as the control.

Table A8 shows that the ordering of the estimates runs against the depth argument. The effect on the aggregate index is largest in the no-prior-claim group, at 0.1175 with a standard error of 0.0376, against 0.0768 (0.0285) for the direct-handover group and 0.0454 (0.0252) for the delayed-handover group. The two handover groups cannot be told apart, and the delayed-handover estimate rests on two districts by the fourth lag. No pairwise difference is individually significant, so what the results support is an ordering. The ordering holds when I pool the two handover groups, which gives 0.0551 against 0.1175, and it holds under the four alternative specifications in Appendix Table A7. The no-prior-claim group's effect is

Table A8: Effect on the aggregate index, by the family's prior hold on the office

Group	Estimate	Joint placebo p-value	Switchers h_1 / h_4
Direct handover: family held it immediately before	0.0768***	0.282	41 / 13
	(0.0285)		
Delayed handover: family held it, a term intervened	0.0454*	0.272	24 / 2
	(0.0252)		
No prior claim: family never held this office	0.1175***	0.357	22 / 10
	(0.0376)		

Each row is a separate estimation on the aggregate index, with that group as the treatment and every other district as the control. Entries are the average cumulative effect per treated district. Estimates come from the heterogeneity-robust estimator of De Chaisemartin and d'Haultfoeuille (2020), on the district-by-wave panel, with districts weighted by their village count. Standard errors, in parentheses beneath each estimate, are clustered on district. Significance: *** 1%, ** 5%, * 10%. The joint placebo column is a nullity test on the lead coefficients. Switcher counts are at the first and fourth lags.

concentrated in household infrastructure, at 0.1416 (0.0438), and it rises to 0.1541 (0.0317) when I drop the districts whose treated elections fall in different groups (Appendix Table A5). Finally, the classification errs in one direction. If I fail to match a relative to the roster, I place a family that held the office in the no-prior-claim group, while the opposite error requires a false match. If a misclassified election behaves like the group it actually belongs to, the gap I measure is smaller than the true gap. The women in the no-prior-claim group come from the broadest networks. Table A9 shows that they have an average of 4.38 office-holding relatives, against 2.20 for the direct-handover group, and 2.86 relatives in legislative office, against 0.80. They also have more relatives at the provincial level, at 0.55 against 0.12. The legislative count is the most reliable row, because the sources name the branch of government in 2,120 of the 2,121 office records. The rows by tier are lower bounds, because 38.7 percent of the records give the branch without the tier.

G.2. The district budget

If the breadth of the network gives a female dynast access to resources from outside the district, it may show up in how her district spends its budget. I estimate the effect on the share of total expenditure that goes to capital, personnel, goods and services, and all other spending, using the annual district budget panel described in Section 4.4. Table A10 shows that the no-prior-claim group is the only group whose budget moves. Its districts shift 2.9 percentage points of total expenditure into capital, which is about twelve percent of the base capital share of 0.235, while personnel falls by 2.4 points and goods and services by 1.8 points.

Table A9: Kinship network breadth by group, means per treated election

	Direct handover	Delayed handover	No prior claim
Office-holding kin	2.20	2.70	4.38
of whom in executive office	1.56	1.53	2.31
of whom in legislative office	0.80	1.40	2.86
Any kin outside the district	0.29	0.20	0.66
Kin at provincial level	0.12	0.10	0.55
Kin at national level	0.02	0.03	0.07

Entries are means per treated election, not estimates. Every measure is a lower bound, because the underlying records give a relative's first office and not her whole career.

The two handover groups show nothing comparable. In logs, capital spending rises by 0.174 (0.075) and personnel is flat at −0.004 (0.033), so the shift reflects more capital spending and not a smaller wage bill (Appendix Table A6).

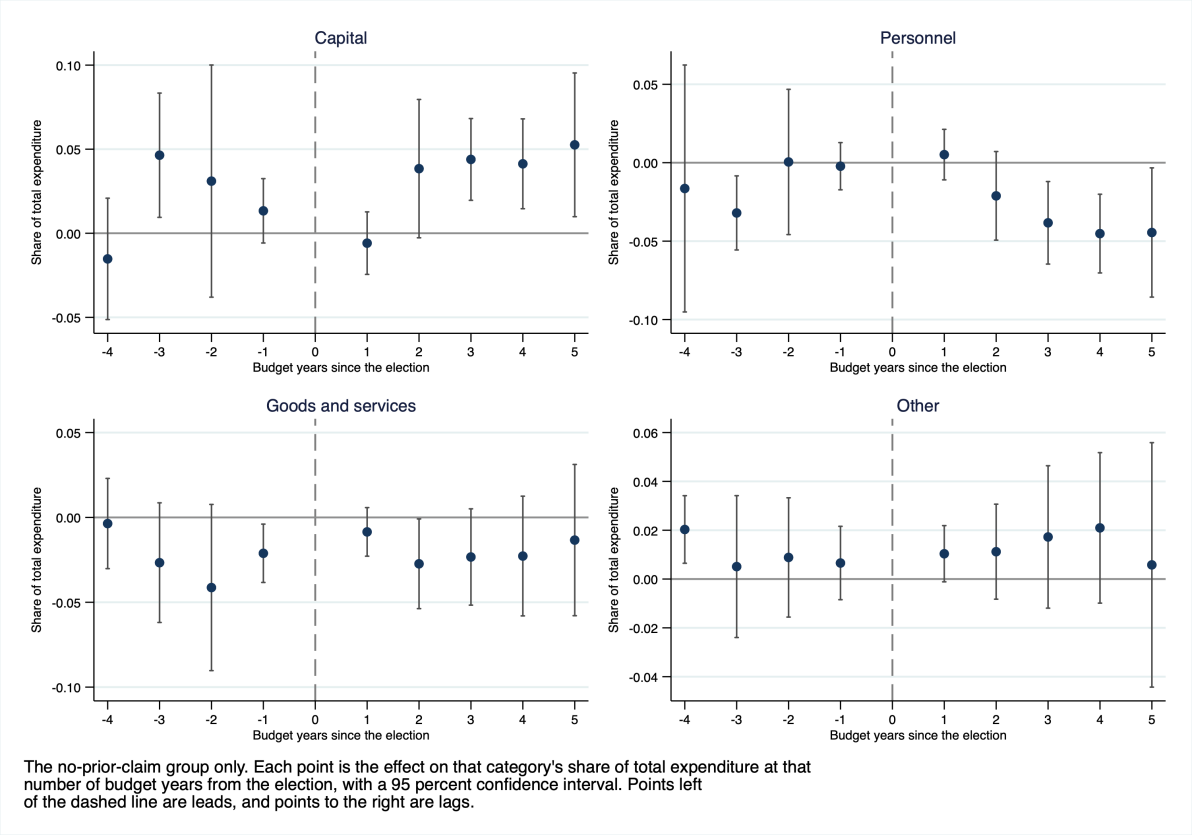
Table A10: Budget composition by the family's prior hold on the office

Group	Capital	Personnel	Goods & services	Other	Sum
All treated	0.0115	-0.0027	-0.0079	-0.0013	-0.0004
	(0.0090)	(0.0080)	(0.0050)	(0.0074)	
Direct handover	0.0043	-0.0034	0.0008	-0.0025	-0.0008
	(0.0110)	(0.0100)	(0.0077)	(0.0099)	
Delayed handover	0.0045	0.0142	-0.0109	-0.0076	0.0002
	(0.0195)	(0.0205)	(0.0106)	(0.0109)	
No prior claim	0.0287***	-0.0237***	-0.0180*	0.0128	-0.0003
	(0.0098)	(0.0079)	(0.0106)	(0.0108)	

Each cell is the average cumulative effect from a separate estimation on the annual district budget panel, 2003–2020. Each outcome is a spending category's share of total expenditure, so 0.01 is one percentage point. The four shares partition total expenditure, so each row sums to approximately zero by construction. Standard errors, in parentheses beneath each estimate, are clustered on district. Significance: *** 1%, ** 5%, * 10%.

There are two limits to this evidence. First, the budget does not show where the additional capital spending comes from. Total revenue and the special allocation transfer (DAK) both fail their placebo tests, so the data cannot confirm that the money comes from a higher level of government (Appendix H). Second, Figure A4 shows that capital and personnel both have a significant lead at $t-3$, in opposite directions and of roughly the same size as the movement after the election, so the shift may have begun before the election. The static two-way fixed effects estimator also finds no effect on capital for this group. The budget results are therefore consistent with the breadth of the network, but they are only suggestive.

Figure A4: No prior claim, event study on shares of total expenditure



Horizon-by-horizon estimates for the no-prior-claim group on each category's share of total expenditure, estimated on the annual district budget panel. Points left of the dashed line are leads and points to the right of it are lags. Bars are 95 percent confidence intervals, clustered on district.

H. Budget results that do not survive their placebos

Two budget results for the no-prior-claim group do not survive their placebo tests, and I set them out here because they bound what the capital result in Appendix G.2 can mean.

The first is total revenue. The placebo coefficients on log total revenue run +0.095, +0.079, +0.099 and +0.032, against effects of +0.030, +0.084, +0.064, +0.059 and +0.060. The placebos are the same size as the effects throughout, so I drop the finding.

The second is the special allocation transfer, DAK. Its share moves for the no-prior-claim group at +0.0142 with a standard error of 0.0057, but in levels log DAK is +0.089 with a standard error of 0.101. The general allocation transfer, DAU, behaves the same way: its share moves at −0.0317 while DAU in levels is flat. Both share results are consistent with an artefact of testing a ratio whose denominator is moving, so I cannot claim that the no-prior-claim group districts bring in more central money.

One related question is settled. The 2016 restructuring of DAK does not produce the

capital result: the interaction with a post-2016 indicator is +0.0005 with a standard error of 0.0078, and the pre-2016 period on its own gives the same magnitude as the full period.

The static two-way fixed effects estimator gives nothing on capital for the no-prior-claim group, at +0.0007 with a standard error of 0.0097. The two estimators do agree on the DAK share, so the disagreement is specific to capital.