Over-Drilling: Local Externalities and the Social Cost of Electricity Subsidies for Groundwater Pumping*

Jesús Bueren — EUI
Xavier Giné — World Bank
Hanan G. Jacoby — World Bank
Pedro Mira — CEMFI

August 2025

Abstract

Borewells for groundwater extraction have proliferated across South Asia, encouraged by subsidized electricity for pumping. Because borewells operating near one another experience mutually attenuated discharges and higher failure rates, farmers deciding whether and when to drill interact strategically with potentially many neighbors through the spatial network of agricultural plots. To incorporate such interactions in policy counterfactuals, we estimate a dynamic discrete network game of well-drilling using plot-level panel data from two states of southern India. We then compare the current regime of free (but rationed) electricity against an annual tax on all functioning borewells that fully defrays electricity costs. We find that the cost-recovery tax, by reining in over-drilling, eliminates a deadweight loss of 135 US$ per acre of land with groundwater potential, 23% of the fiscal cost of the subsidy; nearly half of this social loss can be attributed to the externality. Further, the socially optimal tax rate, which also addresses the local negative externality, is 24% higher than the annual electricity costs. Our estimates also suggest a practical compensation scheme to build farmer support for electricity price reform.

JEL codes: C57, Q15, H23
Keywords: Dynamic network games, Common property resources, Irreversible investment

---

*Bueren email: jesus.bueren@eui.eu, Giné email: xgine@worldbank.org, Jacoby email: hgjacoby@gmail.com, Mira email: mira@cemfi.es. We thank the World Bank Research Committee for financial support, Michael Mann and Siobhan Murray for digitizing the cadastral maps, and Gabriel Aguirre and Bernardo Ribeiro for excellent research assistance. The views expressed in this paper are those of the authors and should not be attributed to the World Bank, its Executive Directors or to the governments they represent. Pedro Mira acknowledges funding from the Spanish Ministry of Science and Innovation (grant PGC2018- 097598-B-I00).

1 Introduction

Private investment in land often affects the returns to investing in neighboring land, but less so for more distant land. Such localized externalities induce strategic interactions, where each landowner's irreversible investment depends on past and expected future investments of their neighbors, who are in turn influenced by the decisions of their neighbors, and so on. This network-game structure presents a significant methodological challenge in quantifying the aggregate welfare implications of a policy change. This paper develops a tractable approach to estimating policy counterfactuals for the case of borewell investment in southern India, a context in which both local spillovers and government intervention in private investment decisions are particularly salient.

Groundwater has become the dominant source of irrigation in India, driving increased agricultural intensification (Jain et al., 2021) and rising rural income (Sekhri 2014). To extract this resource, millions of borewells have sprung up in recent decades, most equipped with submersible electric pumps (Shah 2010; Jacoby 2017). As a borewell operates, the water table around it drops, thus creating a conical draw-down region centered on the pump. If two borewells are close enough to one other, their respective cones of depression overlap, reducing the water flow from each well and possibly contributing to well failure, i.e., when discharge is too low to warrant any cultivation whatsoever.[1] Well interference becomes especially salient in the dry season, during which groundwater is typically the sole source of irrigation. In our study areas — two drought-prone districts in the southern Indian states of Andhra Pradesh and Telengana — these local interference effects dominate the more spatially and temporally diffuse aquifer-level externalities commonly associated with groundwater extraction.[2]

Most Indian states provide farmers free or highly subsidized electricity to run their pumps, artificially inflating the economic returns to well-drilling.[3] Groundwater development has thus devolved into "drilling for subsidies", a form of rent-seeking in which smallholders sink costly wells that would not be economically viable absent these policy-

[1]Blakeslee et al. (2020) find large negative economic impacts of well failure in the South Indian state of Karnataka.

[2]See Sears and Lawell (2018) for a recent review of the large economics literature surrounding the two principal common pool externalities in the context of groundwater, namely: (1) the pumping cost externality whereby one user's pumping, by lowering the watertable in the aquifer, increases the pumping cost of other users, and (2) the strategic (or stock) externality, whereby any groundwater not extracted by a user today may not be available to extract by that user in future years. These externalities are conceptually linked inasmuch as they depend on groundwater being an exhaustible resource. As we discuss in Section 2, groundwater is largely a renewable resource in our context.

[3]In 2013, Indian state governments spent US$11.4 billion to subsidize agricultural power, although this figure likely understates the fiscal drain (Sidhu et al., 2020). Since metering of usage is rare, subsidies generally take the form of low or nonexistent flat charges.

induced distortions (Badiani-Magnusson and Jessoe, 2018). By driving the gross return for the marginal borewell below the private cost of drilling plus the fiscal cost of the subsidy, government policy exacerbates the welfare losses already created by well interference.

Given local externalities, evaluating a policy reform, such as electricity cost-recovery, requires incorporating strategic interactions between neighboring farmers arrayed in a spatial network. Building on the literature surveyed in Aguirregabiria and Mira (2010), we thus develop and estimate a dynamic discrete borewell investment game played on a large (but necessarily bounded) map representing the network of adjacent agricultural plots in the locality. Structural estimation requires taking into account each plot owner's beliefs about the temporal evolution of borewells on all relevant plots, a potentially vast state-space. To avoid this "curse of dimensionality", we make the bounded rationality assumption that investment decisions depend only on beliefs about borewells located in *adjacent* plots. This allows for a novel and tractable estimation strategy that we take to panel data collected in our southern Indian setting. Given parameters of the production technology and other model primitives,[4] we first simulate investment on the plot network map for many periods until a steady state is reached and then compute beliefs based on the temporal evolution of wells in each plot owner's *adjacency*, i.e., the collection of bordering plots determining the local externality. We nest the full solution to this adjacency equilibrium within a Simulated Method of Moments (SMM) estimation algorithm, matching the observed aggregate annual drilling rates by plot size and by the number of currently functional borewells on the plot to their model-based counterparts. As a validation, we use data simulated from the model to closely replicate the reduced-form impact of neighboring borewells on the propensity to drill, a moment not explicitly targeted in our estimation.

Based on the structural estimates, the equilibrium social value of land with groundwater potential is only around 30% of its private value; 70% is accounted for by the capitalized value of the electricity subsidy. If, counterfactually, borewell owners had to make annual payments to fully cover the cost of electricity for pumping, borewell density would fall by 52% and the social value of land would rise by 8,800 Rs/acre (US$ 135). The deadweight loss of free electricity thus amounts to 23% of the fiscal cost of the subsidy. To quantify how the local externality exacerbates this inefficiency, we construct an alternative economy with no well interference and find that deadweight loss in that counterfactual world is only 12% of fiscal cost. Further, we investigate the socially optimal Pigouvian

[4]Specifically, we obtain well interference effects in a first-stage by jointly estimating the probability of well failure and the probabilities of different well flow states as functions of the number of neighboring (functioning) borewells.

borewell tax. While this tax exceeds the cost of electricity by 24%, social welfare is not much higher than in the cost-recovery equilibrium. In other words, once electricity is correctly priced, borewell density decreases to the extent that well interference becomes economically unimportant at the margin. Lastly, we examine the distributional implications of electricity cost recovery, and borewell taxation more generally, suggesting a simple compensation scheme to reduce inequities and build political support for reform.

Despite the prevalence of local investment spillovers, especially in natural resource economics, the literature on dynamic network games remains sparse. To our knowledge, the only other attempt to estimate such a model is Hodgson (2024) in the context of oil field exploration, where the externality is informational and the key inefficiency is thus one of free-riding rather than over-investment.[5] All other empirical applications of network games are static (see, e.g., Acemoglu et al. 2015, Xu 2018, König et al. 2017). In our setting, a static model in which all plot owners sink their borewells at once would not account for the empirically important and inherently dynamic feature of well failure. A static model also cannot speak to the short-run distributional implications of a borewell tax. By contrast, our dynamic model allows us to calculate how such a policy change impacts landowners with and without functioning borewells at baseline along the entire transition path to the new steady state, thereby informing the design of practical compensation schemes. More fundamentally, welfare analyses of policies affecting irreversible investments, such as borewells, call for both a dynamic model *and* the consideration of transition paths, as comparisons of long-run steady states can be misleading (Domeij and Heathcote, 2004). Finally, a dynamic model enables us to exploit data on drilling choices conditional on the current number of functioning borewells – i.e., the dynamic decision rule – for identification of structural parameters.

As the first to incorporate well interference externalities in a full dynamic general equilibrium model of drilling decisions, this paper also contributes to the economics of groundwater extraction, a literature focused almost exclusively on the United States and, to a lesser extent, on the world's largest groundwater user, India. Pfeiffer and Lin (2012) considers well interference externalities in the High Plains Aquifer of Kansas whereas Sears et al. (2022a) studies strategic interactions among neighboring extractors in California. Both of these papers are concerned with the intensive margin (pumping) as opposed to the extensive margin (drilling) and neither models the spatial network game. Groundwater research in India focuses on the the deep alluvial aquifers of the northwest, where 'mining' of fossil groundwater is a serious concern (see Fishman et al. 2011). In this context,

[5]The pioneering work of Lin (2013) in the same context is limited to studying strategic interactions between isolated adjacent pairs of oil parcels thus abstracting from the spatial network problem.

the calibration study of Sayre and Taraz (2019) combines decisions about both groundwater pumping and investment in deeper wells within a partial equilibrium framework, i.e., ignoring well interference externalities. Ryan and Sudarshan (2022) econometrically estimates the welfare cost associated with over-pumping in Rajasthan. By focusing on the aquifer-wide pumping cost externality (see footnote 2) and by taking the number of borewells as fixed, their paper abstracts from both well interference and from drilling costs. Ryan and Sudarshan (2022) finds that electricity rationing to agriculture leads to roughly the socially optimal quantity of groundwater pumping on average. Similar rationing in our study areas also limits over-pumping, so that whatever intensive margin distortion remains is likely to be small in comparison to the extensive margin distortion that we emphasize.

More broadly, this paper contributes to a growing literature estimating the value of water in the context of missing markets. A non-exhaustive listing of recent work includes: Rafey (2023), Donna and Espín-Sánchez (2023), and Hagerty (2023) on surface irrigation; Carleton et al. (2023) and Sekhri (2022) on the role of international trade; and Burlig et al. (2021) and Bruno and Jessoe (2021) on the price-elasticity of groundwater demand.

The next section of the paper describes the setting and data. Section 3 lays out the model of borewell investment and local externalities. Section 4 discusses the structural estimation procedure and presents the results. Section 5 considers counterfactual policies, including the optimal borewell tax, and Section 6 concludes the paper.

2 Setting and Data

2.1 Context

Before its partition into Andhra Pradesh (AP) and Telangana in 2014, unified Andhra Pradesh was one of the most important agricultural states of India, accounting for about 7 percent of gross cropped area nationally, roughly half of which was irrigated with groundwater. However, as argued by Kumar et al. (2011), the economic efficiency of groundwater extraction in unified AP has decreased substantially. A tripling in the number of borewells from 1995 to 2010 to more than 1.5 million (Jacoby 2017) has led to high rates of well failure, lower irrigated area per well, and higher energy requirements for groundwater pumping due to well interference. Meanwhile, power supply to agriculture for running electrical pumps has become a political issue all over India. In 2004, a newly elected government of unified AP abolished flat rate electricity charges, which had covered just 11 percent of the provision cost, making agricultural power free to farmers, a move swiftly followed by the major states of Tamil Nadu, Karnataka, and Punjab.[6] Currently, farmers in AP and Telangana typically run their pumps continuously during the 7 to 9 hours per day when this free electricity is made available for agricultural use (Fishman et al., 2023).

Figure 1: WATER TABLE FLUCTUATIONS: 1998-2014

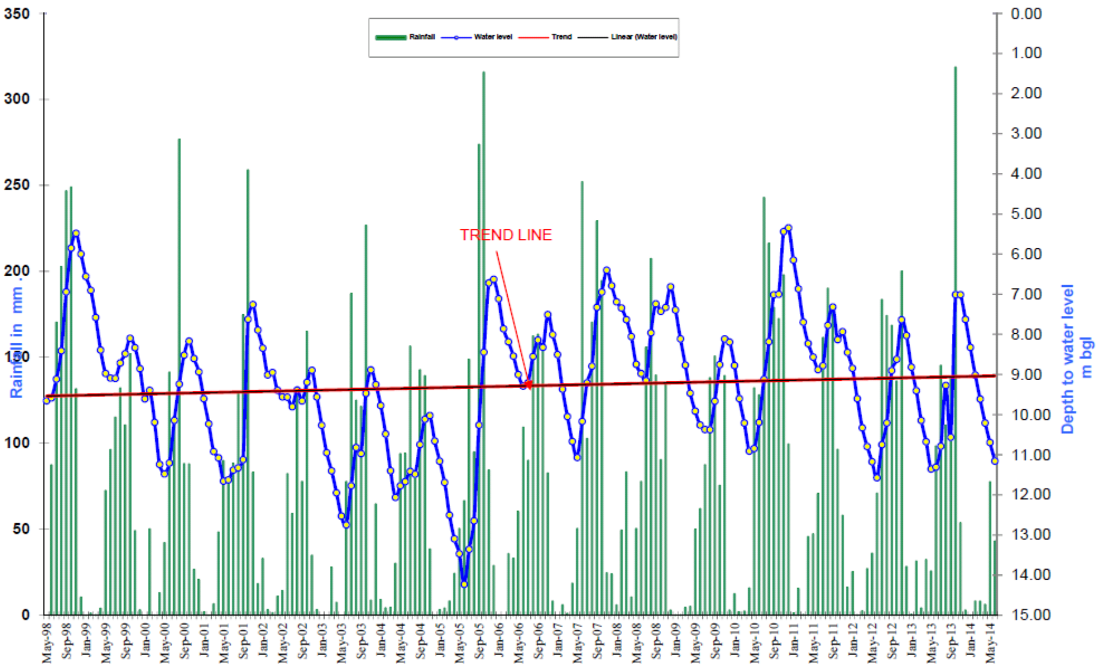

Notes: Average depth to water table in meters below ground-level from all state observation wells and rainfall in millimeters by month (Source: Andhra Pradesh Groundwater Department, http://apsgwd.gov.in/swfFiles/reports/state/monitoring.pdf; last accessed Feb. 10, 2016).

In contrast to northwest India, where deep groundwater reserves are available to mine, much of south India is underlain by shallow hard rock aquifers with limited groundwater storage capacity. Recharge from monsoon rains is mostly depleted through pumping during the subsequent dry season. Figure 1 indicates that the time-series of depth to water table across unified AP, a measure of overall resource depletion, is dominated by intra-annual variability, showing practically zero trend from 1998-2014, the most recent years for which we have consistent data before the partition.[7] Because of their low transmissivity

[6]Shah et al. (2012) estimates that these subsidies in AP amounted to 94% of the gross value of its agricultural output before partition. The corresponding figure in the more agriculturally productive state of Punjab is only 12%. Note that Shah et al. (2012) uses an annual electricity cost per borewell of about US$450 for the entire state of AP circa 2010, whereas we obtain a much more conservative figure of US$180 (8,500 Rs) in our study areas (see Appendix A).

[7]Hora et al. (2019) argues that such water table trends are biased upward by relying on surviving (i.e., non-failed) observation wells to measure groundwater levels across time. Indeed, our analysis of well failure in Appendix F is consistent with a secular, but rather slow, decline in water tables in our study area.

(velocity of horizontal groundwater flow), hard-rock aquifers accentuate well interference. In our setting, an interwell spacing of at least 250 meters is recommended to avoid interference effects (Chandrakanth 2015). Blakeslee et al. (2020) detail the process of well failure in hard-rock aquifers, highlighting *local* hydrogeological features, i.e., sub-surface fractures fed from different sources of recharge, rather than aquifer-wide depletion. In our setting, therefore, well interference is the predominant groundwater pumping externality, one that is both localized and static, affecting only current groundwater availability.

2.2 Household surveys and representative plot sample

Our data come from the drought-prone districts of Anantapur (Andhra Pradesh) and Mahabubnagar (Telangana), originally the backdrop for the weather insurance study of Cole et al. (2013). As shown in Giné and Jacoby (2020), groundwater availability and the related development of groundwater markets in these drought-prone districts is extremely limited compared to districts that receive the average amount of annual rainfall and, especially, to those in water-abundant coastal AP. Only farmers with access to a functioning borewell can cultivate during the dry (rabi) season, when they mainly grow paddy, groundnut, maize, and mulberry in Anantapur and paddy and groundnut in Mahabubnagar. In the wet (kharif) season, farmers in both districts grow primarily paddy, sorghum, and groundnut and use groundwater to supplement monsoonal rainfall.

Highly fragmented landownership also contributes to well interference externalities. To obtain the typical spatial layout of separately owned plots, we digitized cadastral maps for at least one village in each of the 12 mandals (sub-districts or counties) covered by our survey (see Appendix B). In all, we have 14 such village maps containing 12,330 land parcels; the median plot size is only 2.02 acres.

Representative plot sample In 2017, we were able to re-interview 1,436 of 1,488 randomly selected farm (landowner) households originally surveyed in 2010 by Cole et al. (2013). The 2017 survey includes a history of well-drilling attempts on and around each of the household's plots since 2011 and records every borewell present on each plot regardless of whether currently functional (i.e., has non-negligible discharge) or even dismantled. From this information, we construct a retrospective five-year panel of drilling attempts and the number of functional borewells on 2,862 plots, referring to this as the representative plot sample. Highlighting its representativeness, the median plot area of 2.00 acres in this sample virtually coincides with the median area of plots in the 14 digitized cadastral maps mentioned above. The 2010 and 2017 surveys also collected data on discharge or flow for every functioning borewell for the 2009-10 and 2016-17 rabi seasons (see Appendix F).

Timeline Figure 2 provides an event timeline to guide our empirical and theoretical analyses. A "year" begins with the onset of the monsoon (initiating the kharif or wet season), followed by rabi (dry) season planting once the rains have ceased. Borewells are drilled in the pre-monsoon (summer) season when water tables are at their lowest; this assures farmers that, if successful, the new borewell will yield groundwater throughout the rabi season. New borewells are thus available for pumping only in the year following a successful drilling attempt, with year t "success" defined as being functional at least during year $t+1$. Consistent our survey's elicitation of whether a borewell is functional or not, we assume that failures are *realized* during rabi season planting. For the plot-level retrospective panel, we drop data from 2017 because, as the survey was administered in May, not all drilling attempts or well failures that occurred in 2017 were necessarily captured in the dataset.[8]

Figure 2: Timeline

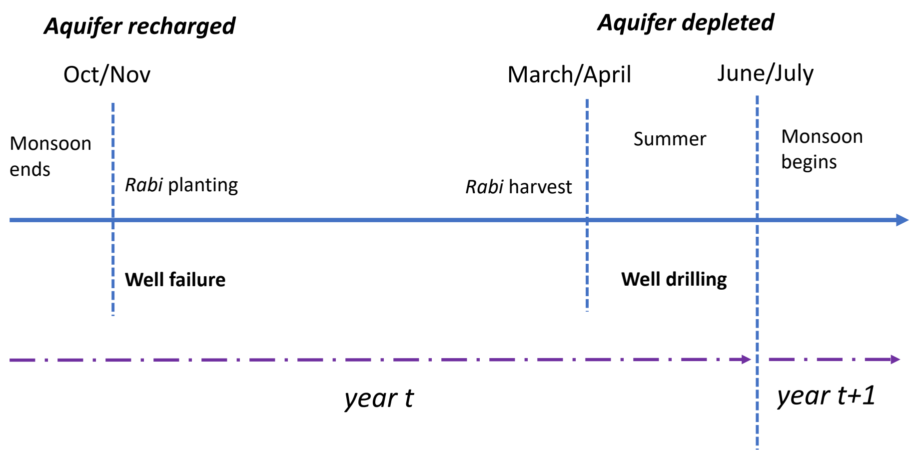

Drilling patterns There were 526 attempts to drill a borewell made on 437 plots between 2012-16, only 197, or 37.5%, of which were successful (i.e., resulted in a functional borewell).[9] Panel (a) of Figure 3 shows annual drilling rates by the number of functioning borewells and plot area quintiles. While the propensity to drill is *higher* on plots

[8]For consistency with the adjacency survey panel (see below), we also drop data from 2011.

[9]Multiple annual drilling attempts occurred in 134 (25%) of these 526 cases (with a mean of 2.5 attempts conditional on a multiple attempt). Since most of these multiple attempts likely entail farmers drilling on different portions of their plot in search of a good water source, rather than attempting to sequentially drill multiple successful wells in one year, we code them all as a single 'attempt' (and,

Figure 3: Drilling attempts

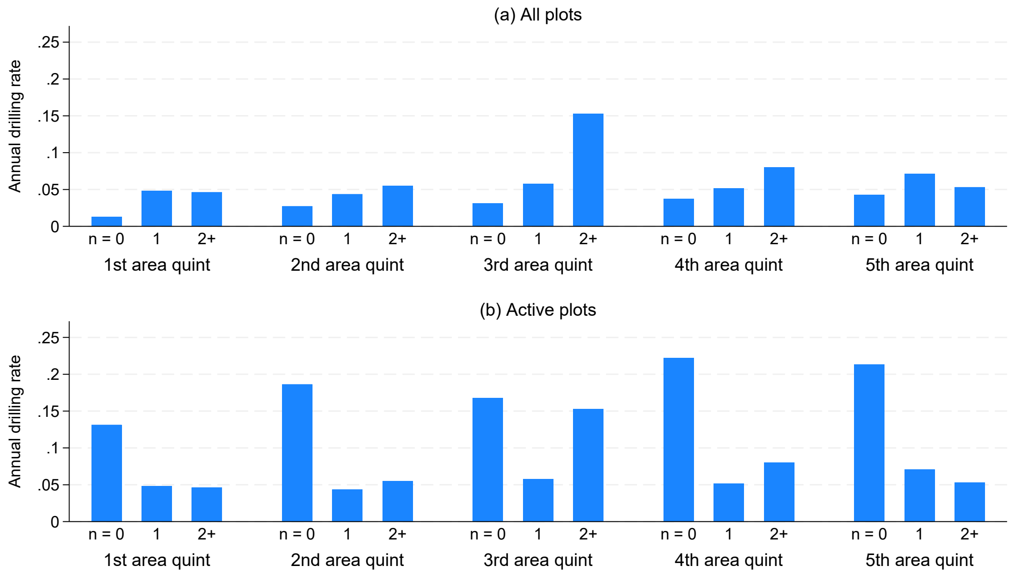

Notes: Average annual drilling attempt rate by plot area quintiles and by number (n) of functioning borewells on the plot. Panel (a) is for the representative sample of plots; panel (b) is for active plots only (i.e., those which had any functioning borewell or drilling attempt during 2012-16).

with more functioning borewells, this may reflect heterogeneity in groundwater potential: where drilling is profitable, there is both more drilling *and* more functioning borewells. Panel (b) restricts attention to plots in which at least one drilling attempt was made during 2012-16 or that already had at least one functioning borewell in 2012. The owners of such "active" plots (39% of the representative sample) evidently believed that they had potential for groundwater development. The "active" plot sample shows a negative relationship between functioning borewells and drilling, consistent with diminishing marginal returns to investment. Thus, conditional drilling moments will provide identifying information about the underlying production function. The contrast between panels (a) and (b) further indicates that only a subset of plots (imperfectly proxied by their "active" status) may be suitable for groundwater development. We later term these suitable plots "developable" and treat this unobserved type ($\mathcal{D}$) as a latent factor in our structural estimation.

correspondingly, in the model only allow one borewell to be drilled each year). As support for this choice, we note that, in the representative panel of plots, only 7 out of 11,448 (0.06%) year-to-year changes in functioning well numbers are from zero to two (or more).

Adjacency survey We define an *adjacency* as the set of all agricultural plots contiguous to a reference plot, inclusive of it. As part of the 2017 household survey, an adjacency survey was administered covering 1,057 farmers with an eligible reference plot. Eligibility required that at least one drilling attempt had been made in the last seven years either on the plot itself or within a 500 meters radius of the plot. (If the household had two or more eligible plots, one was chosen at random). The adjacency survey asks each reference plot owner for retrospective information annually, going back to 2011, about the existence and status (functioning or not) of all borewells in the adjacency. Following our timing conventions, we match drilling activity and borewell failure on reference plot i in year t with the number of functioning wells on the reference plot, n_{it}, and with the number of functioning wells in the adjacency outside of reference plot, $\mathcal{N}_{it}$, both observed at the *beginning* of year t, i.e., before any year t failures. (We drop data from 2011 because we do not have $\mathcal{N}_{it}$ for that year).

Throughout the paper, we denote the total number of functioning wells in the adjacency by N_{it}, where $N_{it} \equiv \mathcal{N}_{it} + n_{it}$. Since adjacency survey respondents may recall the functioning status of borewells on other farmers' plots less accurately than of those on their own (reference) plots, our econometric work based on adjacency survey data allows for $\mathcal{N}_{it}$ (but not n_{it}) to be measured with error, i.e., misclassified as functioning when actually failed or as failed when actually functioning.[10] We assume perfect recall for all other measurements in the adjacency and household surveys.

3 A Model of Borewell Investment

In this section, we present a dynamic equilibrium model of borewell investment on a large plot network. Based on evidence adduced in Appendix C, we abstract from liquidity constraints and other financial frictions.[11]

[10] Specifically, our estimates of the well failure process (Section 4.3 and Appendix F) and the model validation exercise (Section 4.5 and Appendix H) rely on adjacency survey data. To deal with misclassification in $\mathcal{N}_{it}$, we use the number of *existing* wells in year t as an upper bound for the unobserved true number of functional wells in the borewell failure model, or as an instrument for the reported number of functional wells in the model validation exercise. We argue that existing wells, i.e., those that were ever functional, are reported accurately, even if their functional status and the exact timing of failure events, is subject to recall error.

[11] In particular, we find no link between the pre-sample wealth of the plot owner and their propensity to drill from 2012-16, and we show descriptive evidence of farmers' access to credit.

3.1 Preliminaries

Let the *incremental* output from groundwater irrigated agriculture on plot i at time t be

$$y_{it} = \theta \left[\alpha (\sum_{w} q_{it}^{w})^{\delta} + (1-\alpha) a_i^{\delta} \right]^{\frac{1}{\delta}}, \tag{1}$$

where (θ, α, δ) are parameters, a_i is plot area, and q_{it}^{w} is discharge or flow from well w. Yearly flow is stochastic and thus unknown to the farmer prior to drilling and has a discrete distribution with K points of support $\{q_{it1}^{w}, ..., q_{itK}^{w}\}$ each with probability π_{itk}. Along with constant elasticity of substitution (CES), the production function defined in (1) imposes constant returns to scale (CRS); i.e., output per acre depends only on flow per acre.[12] The scale parameter θ converts physical output into 2017 Indian rupees (Rs).

We assume that only functioning borewells within plot i's adjacency influence flow and failure of borewells on that plot; borewells outside of the adjacency cause no interference. This is a reasonable assumption given the typical size of plots and range of well interference effects in our setting.[13] Year t flow-state probabilities π_{itk} thus depend on the number of functioning borewells in the adjacency at the beginning of year t, as well as on that year's monsoon rainfall (i.e., aquifer recharge), according to

$$\pi_{itk} = \pi_k(N_{it}, R_t). \tag{2}$$

Due to increased well interference, higher N_{it} shifts probability mass to low flow states.

A borewell remains functional, with positive discharge, until stochastic failure occurs at rabi planting with probability $\pi_{Fit} = \pi_F(N_{it-1}, R_t)$, which again incorporates the localized externality. We assume that failure is an absorbing state, so that once a borewell fails it ceases to discharge water for good. Put differently, there is no possibility of 'resuscitation'.[14,15]

[12] The Online Appendix of Giné and Jacoby (2020) tests and cannot reject CRS based on a Cobb-Douglas production function estimation in a closely related setting.

[13] In a chessboard configuration of identical plots averaging one hectare (as in our data) with borewells located at the center, the distance between a borewell in the reference plot and one elsewhere in the adjacency would be 100-140 meters, well within the range of interference effects in Chandrakanth (2015). Expanding the definition of adjacency to include a second ring of identical plots would increase the average distance between wells to 200-280 meters, beyond the range of significant interference effects.

[14] This is a simplification for the sake of empirical tractability; resuscitation does occur, albeit infrequently. Examining failure histories of the 429 borewells in our survey reported to have first failed in the decade between 2005-2014, 48 (or 11%) began to function again up to 3 years later, although 3 of these wells had failed again by 2017. We code wells that resuscitate within 2 years of first failure (most cases) as always functional and wells that resuscitate only after 3 years as failed in the first year of failure.

[15] While we allow the failure probability to depend on rainfall for the sake of generality, we find a null

We further assume that the probability of a successful drilling attempt, π_S, is constant. Each attempt to drill a borehole costs c_d and, if successful, entails an additional cost of installing a pipe, casing, and hooking up the electrical connection; the submersible pump itself is removable and thus we do not consider it a sunk cost. The total cost of a successful attempt is, therefore, $c_s > c_d$.

Finally, for the sake of tractability, we assume that at most two wells can function simultaneously on any given plot so that $N_{it} \in \{0, .., 2p_i\}$, where p_i is the number of plots in adjacency i.[16] Drilling success, failure, and discharge events for two wells on the same plot are independent random variables *conditional* on the plot-specific unobserved heterogeneity (described in Section 4). As in equation (1), the incremental output of a plot with two borewells depends on the sum of their discharges $q^1_{it} + q^2_{it}$, since water from both wells can be pooled and dispersed throughout the plot. Expected output conditional on, respectively, the number of functioning borewells in the adjacency, the number in the farmer's own plot, and on monsoon rainfall, $Y_{it}(N_{it}, n_{it}, R_t)$, is thus

$$\begin{aligned} Y_{it}(N_{it}, n_{it}, R_t) &= \sum_{k=1}^{K} \pi_{itk}(N_{it}, R_t)\theta \left[\alpha(q^1_{itk})^\delta + (1-\alpha)a_i^\delta\right]^{\frac{1}{\delta}} && \text{if } n_{it} = 1 \\ &= \sum_{j=1}^{K}\sum_{k=1}^{K} \pi_{itj}(N_{it}, R_t)\pi_{itk}(N_{it}, R_t)\theta \left[\alpha(q^1_{itk} + q^2_{itj})^\delta + (1-\alpha)a_i^\delta\right]^{\frac{1}{\delta}} && \text{if } n_{it} = 2. \end{aligned} \tag{3}$$

3.2 Borewell investment decision

We now consider the discrete choice to drill ($d = 1$) or not to drill ($d = 0$) and derive the plot owner's decision rule or conditional choice probability $CCP(\mathcal{N}, n) \equiv \Pr(d = 1 \mid \mathcal{N}, n)$, temporarily dropping subscripts for ease of exposition. We first describe the dynamic decision facing the owner of a plot with area a in an adjacency with p plots, in isolation, i.e., taking as given their beliefs about the evolution of the state of the adjacency. As noted, the state space of the plot owner consists only of the total number of wells in the other plots of the adjacency $\mathcal{N} \in \{0, .., 2(p-1)\}$ and the number of own functioning wells, $n \in \{0, 1, 2\}$. In the next subsection and later in Section 4.1, we discuss this assumption and its role in a tractable equilibrium model of beliefs and conditional choice probabilities.

By assumption, state $n = 0$ or $n = 1$ are the only cases where investment can occur. A

effect of rainfall in Section 4.3, consistent with our assumption of well failure being an absorbing state.

[16]In our representative plot panel, 3 or more functioning wells occurs in just 35 out of all 14,310 plot-years.

plot owner with $n = 0$ may decide not to drill, with payoff value $\overline{v}_{00}(\mathcal{N}) + \epsilon_{00}$, or to drill, with payoff value $\overline{v}_{0I}(\mathcal{N}) + \epsilon_{0I}$. As in a random-utility framework, choice-specific payoffs have additive "deterministic" and "random" components. The random components of the payoff of waiting (ϵ_{00}) or drilling (ϵ_{0I}) are realized every period before choices are made, iid across choices and time, and unobserved by other plot owners in the adjacency, each of whom are drawing their own random components.

The deterministic components, which are known to the plot owner conditional on the observable state variables and parameters, include the static one-period profits (expected value of output minus drilling costs, if any) and the expected continuation values. For the no drilling (waiting) choice,

$$\begin{aligned} \overline{v}_{00}(\mathcal{N}) &= \beta \, \mathbb{E} \, V(\mathcal{N}', 0) \\ &= \beta \sum_{\mathcal{N}'} \tilde{F}(\mathcal{N}' \mid \mathcal{N}, 0) V(\mathcal{N}', 0) \end{aligned} \tag{4}$$

and for the choice of making a drilling attempt

$$\begin{aligned} \overline{v}_{0I}(\mathcal{N}) = \pi_S & \left(-c_s + \beta \sum_{\mathcal{N}'} \tilde{F}(\mathcal{N}' \mid \mathcal{N}, 0) V(\mathcal{N}', 1) \right) \\ & + (1 - \pi_S) \left(-c_d + \beta \sum_{\mathcal{N}'} \tilde{F}(\mathcal{N}' \mid \mathcal{N}, 0) V(\mathcal{N}', 0) \right), \end{aligned} \tag{5}$$

where the value function $V(\mathcal{N}, n)$ is defined below, β is the discount factor and $\tilde{F}(\mathcal{N}' \mid \mathcal{N}, n)$ reflects beliefs about the probability of $\mathcal{N}'$ functioning wells in other adjacency plots next period conditional on $\mathcal{N}$ functioning wells in other adjacency plots and on n functioning wells on the reference plot ($n = 0$, in this case) this period. Since drilling occurs after rabi season production (see Figure 2), the increase in expected output from any successful attempt is only realized in the next period.

We assume that the random components associated with the choices of waiting and drilling, $(\epsilon_{00}, \epsilon_{0I})$, are each iid Type-1 extreme value with location parameter 0 and scale parameter σ. Further, denote by $V(\mathcal{N}, n)$ the beginning-of-period value function for the plot owner, before these random components of payoffs are realized. Taking expectations for $n = 0$, we have

$$\begin{aligned} V(\mathcal{N}, 0) &= \mathbb{E} \max \Big\{ \overline{v}_{00}(\mathcal{N}) + \epsilon_{00}, \overline{v}_{0I}(\mathcal{N}) + \epsilon_{0I} \Big\} \\ &= \sigma \Big(\gamma + \log \big(\exp(\overline{v}_{00}(\mathcal{N})/\sigma) + \exp(\overline{v}_{0I}(\mathcal{N})/\sigma) \big) \Big) \end{aligned} \tag{6}$$

where the second line follows from the Type-1 extreme value assumption and γ is Euler's constant.

A borewell owner with $n = 1$ may decide to wait or attempt a second borewell. Waiting yields payoff value $\overline{v}_{10}(\mathcal{N}) + \epsilon_{10}$, where

$$\begin{aligned}\overline{v}_{10}(\mathcal{N}) =& Y(\mathcal{N}+1,1,R) + \beta\,\mathbb{E}\Big[(1-\pi_F(\mathcal{N}+1,R'))\sum_{\mathcal{N}'}\tilde{F}(\mathcal{N}' \mid \mathcal{N},1)V(\mathcal{N}',1) \\ &+ \pi_F(\mathcal{N}+1,R')\sum_{\mathcal{N}'}\tilde{F}(\mathcal{N}' \mid \mathcal{N},1)V(\mathcal{N}',0)\Big],\end{aligned} \tag{7}$$

with expectations taken with respect to the distribution of next period's monsoon, R'. Equation (7) reflects the timeline in Figure 2, with current year drilling decisions made before the farmer knows about failures of currently functioning borewells; these will only be realized after the upcoming monsoon and once next year's rabi planting starts. Attempting a second borewell yields payoff value $\overline{v}_{1I}(\mathcal{N}) + \epsilon_{1I}$, where

$$\begin{aligned}&\overline{v}_{1I}(\mathcal{N}) = Y(\mathcal{N}+1,1,R) - c_s\pi_S - c_d(1-\pi_S) \\ &+ \beta\,\mathbb{E}\Big[\pi_S(1-\pi_F(\mathcal{N}+1,R'))\sum_{\mathcal{N}'}\tilde{F}(\mathcal{N}' \mid \mathcal{N},1)V(\mathcal{N}',2) \\ &+ (\pi_S\pi_F(\mathcal{N}+1,R') + (1-\pi_S)(1-\pi_F(\mathcal{N}+1,R')))\sum_{\mathcal{N}'}\tilde{F}(\mathcal{N}' \mid \mathcal{N},1)V(\mathcal{N}',1) \\ &+ (1-\pi_S)\pi_F(\mathcal{N}+1,R')\sum_{\mathcal{N}'}F(\mathcal{N}' \mid \mathcal{N},1)V(\mathcal{N}',0)\Big].\end{aligned} \tag{8}$$

We can now write

$$\begin{aligned}V(\mathcal{N},1) &= \mathbb{E}\max\Big\{\overline{v}_{10}(\mathcal{N}) + \epsilon_{10}, \overline{v}_{1I}(\mathcal{N}) + \epsilon_{1I}\Big\} \\ &= \sigma\big(\gamma + \log\left(\exp(\overline{v}_{10}(\mathcal{N})/\sigma) + \exp(\overline{v}_{1I}(\mathcal{N})/\sigma)\right)\big)\end{aligned} \tag{9}$$

where the second line follows, again, from the Type-1 extreme value assumption. Finally, a farmer with two functioning borewells on a plot cannot invest in a new one and, as a result, we have

$$
\begin{aligned}
V(\mathcal{N},2) = Y(\mathcal{N}+2,2,R) + \beta\,\mathbb{E}\Big[&(1-\pi_F(\mathcal{N}+2,R'))^2 \sum_{\mathcal{N}'} \tilde{F}(\mathcal{N}' \mid \mathcal{N},2) V(\mathcal{N}',2) \\
&+ 2\pi_F(\mathcal{N}+2,R')(1-\pi_F(\mathcal{N}+2,R')) \sum_{\mathcal{N}'} \tilde{F}(\mathcal{N}' \mid \mathcal{N},2) V(\mathcal{N}',1) \\
&+ \pi_F^2(\mathcal{N}+2,R') \sum_{\mathcal{N}'} \tilde{F}(\mathcal{N}' \mid \mathcal{N},2) V(\mathcal{N}',0)\Big].
\end{aligned} \tag{10}
$$

Equations (4)-(10) combine to form the Bellman equation for this investment problem.

The discrete choice to attempt drilling a borewell in the reference plot is thus

$$
d = d(\mathcal{N},n) = \begin{cases} 1 & \text{if } n < 2 \text{ and } \overline{v}_{nI}(\mathcal{N}) - \overline{v}_{n0}(\mathcal{N}) > \epsilon_{n0} - \epsilon_{nI} \\ 0 & \text{otherwise,} \end{cases}
$$

using equations (4), (5), (7) and (8). With logit random utility shocks, the decision rule as perceived by the researcher (and by neighbors) is characterized by the CCP function

$$
\begin{aligned}
\mathrm{CCP}(\mathcal{N},n) &= \Pr(d=1 \mid \mathcal{N},n) = \Pr(\epsilon_{n0} - \epsilon_{nI} < \overline{v}_{nI}(\mathcal{N}) - \overline{v}_{n0}(\mathcal{N})) \\
&= \frac{\exp(\overline{v}_{nI}(\mathcal{N})/\sigma)}{\exp(\overline{v}_{nI}(\mathcal{N})/\sigma) + \exp(\overline{v}_{n0}(\mathcal{N})/\sigma)}.
\end{aligned}
$$

3.3 Adjacency equilibrium

Before characterizing the equilibrium of the dynamic drilling game, we introduce the concept of a village "map", or plot network, upon which this game is played. Although the borders of our administrative maps (see Section 2.2 and Appendix B) do not generally correspond to salient geographic or geological features, each map contains many plots and as a result "truncation-at-border" effects should have negligible empirical consequences.[17] Formally, a cadastral map with P plots is characterized by a $P \times 1$ vector A listing the area of each plot and a $P \times P$ adjacency matrix $\mathbf{M}$ with typical element $M_{ij} = 1$ if plot j adjoins plot i and 0 otherwise, and with $M_{ii} = 1$. Ignoring for now plot-specific heterogeneity, $\{A, \mathbf{M}\}$ fully characterizes all adjacencies in the map. For instance, plot i has an area equal to the i-th element of A and its adjacency has $\sum_j M_{ij}$ plots because plot

[17] Absent natural spatial "breaks" in well interference externalities, a tractable empirical model requires a bounded plot network, which necessarily generates truncated adjacencies; plots at the map-village border have at least one adjacent plot that lies in a different cadastral map (unavailable to us). Truncation induces an error in computing the adjacency equilibrium (discussed below) because, strictly speaking, the true equilibrium is determined by interference effects from plots outside the cadastral map, which we ignore. Arguably, however, this error is small because our average cadastral map has 881 plots with only 102 (12%) being border plots.

j with $M_{ij} = 1$ belongs in plot i's adjacency. Let $\mathcal{M}_{(ih)}$ be the set of plots h-level adjacent to plot i so that $\mathcal{M}_{(i1)} = \{j : M_{ij} = 1\}$ is the set of immediate (1-level) neighbors in i's adjacency, $\mathcal{M}_{(i2)} = \left\{j : j \notin \mathcal{M}_{(i1)}, \exists k : j \in \mathcal{M}_{(k1)}, k \in \mathcal{M}_{(i1)}\right\}$ is the set of 1-level adjacent neighbors of i's 1-level adjacent neighbors, and so on for all "layers" h.

Let the state of plot i in period t be the number of functioning wells on the plot at the beginning of the year $n_{it} \in \{0, 1, 2\}$. Further, let $X_t = \{n_{it} : i = 1, \ldots, P\}$ be the state of the map, representing the entire spatial distribution of borewells in the cadastral map. Now, define $X_{(ih)t} = \left\{n_{jt} : j \in \mathcal{M}_{(ih)}\right\}$, where $X_{(i1)t}$ collects the state of the neighbors of reference plot i, $X_{(i2)t}$ collects the state of the neighbors' neighbors, and so on.

Thus far, we have taken beliefs about the evolution of the number of functioning wells in the adjacency as given, viewing the plot owner's investment decision as a game "against nature". Consider now a Markov-perfect equilibrium (MPE), in which beliefs and decision rules (CCPs) of all plot owners are consistent with one another. Our state space $(\mathcal{N}, n)$ implicitly assumes that plot owners ignore the status of wells on successive layers of plots outside their own adjacencies. This restriction is not, in general, implied by our key assumption that well interference is limited to functioning wells in the adjacency. Indeed, information on the status of wells in the second layer might help agents predict immediate neighbors' investment behavior and thus the status of wells in the adjacency, information about the third layer might help predict the status of wells in the second layer, and so on. Under *unrestricted* MPE play, therefore, investment decisions may depend on the state of the whole map, even with well interference effects confined to adjacent plots.

To be precise, let $CCP_i(X_t)$ be a choice probability function for the owner of plot i and $\{CCP\}$ be the vector of choice probabilities of all plot owners in the cadastral map. Further, let one-period ahead transition probabilities $\widetilde{F}(X_{t+1} \mid X_t)$ describe beliefs about the evolution of the state of the map and $F(X_{t+1} \mid X_t; \{CCP\})$ be the one-period-ahead law of motion for the state induced by the primitives and $\{CCP\}$. We thus have:

Definition 1. An MPE is a vector of choice probabilities $\{CCP_i^*(X_t) : i = 1, \ldots, P\}$ and beliefs $\widetilde{F}^*$ such that: a) given $\widetilde{F}^*$, CCP_i^* is the solution of plot owner i's dynamic game "against nature"; and b) beliefs are correct, in that $\widetilde{F}^*(X_{t+1} \mid X_t) = F(X_{t+1} \mid X_t; \{CCP\}^*)$.

In general, each plot owner in their unique adjacency would have a different equilibrium CCP depending on all primitives, including the structure of the map. Given the number of plots in the map, unrestricted MPE play is empirically infeasible due to the high dimensionality of $\{X_t, \{CCP\}\}$.

As a tractable alternative, we consider a Markov equilibrium in which: i) CCPs depend only on the state of the (1-level) adjacency $(X_{(i1)t}, n_{it})$, and ii) the plot owner has beliefs only about the stochastic evolution of $X_{(i1)t}$ in steady state. While assumption i) avoids the "curse of dimensionality", the fact that well interference is limited to the adjacency should dampen the influences induced by unrestricted play of plot owners in layers $h > 1$ as well as make it less plausible (i.e., by bounded rationality) that plot owners would keep track of the full state of a large map. Assumption ii), specifically that plot owners *only* have beliefs about the state of their own adjacency, is a natural implication of assumption i) but it also adds the non-trivial simplification that equilibrium beliefs are correct when averaged over the map's stochastic steady state. Thus, in the spirit of an "oblivious equilibrium",[18,19] we propose

Definition 2. An Adjacency Equilibrium (AE) is a vector of choice probabilities $\{CCP_i^*(X_{(i1)t}, n_{it}) : i = 1, \ldots, P\}$ and of beliefs $\left\{\widetilde{F}_i^*(X_{(i1)t+1} \mid X_{(i1)t}, n_{it}) : i = 1, \ldots, P\right\}$ such that: a) given beliefs $\widetilde{F}_i^*$, CCP_i^* is the solution of plot owner i's dynamic game "against nature"; and b) beliefs are correct "on average" in steady state. That is, let $F^\infty(X_t; \{CCP\})$ be the stationary joint distribution over the state induced by the primitives and the vector of CCPs,[20] and let $F_i(X_{(i1)t+1} \mid x_{(i1)t}, x_{(i2)t}; \{CCP\})$ be the one-period-ahead law of motion for x_{i1} induced by the same. Further, let $F_i^\infty(X_{(i2)t} \mid X_{(i1)t}, n_{it}; \{CCP\})$ be the conditional distribution implied by $F^\infty(X_t; \{CCP\})$. Then,

$$\begin{aligned} &\widetilde{F}_i^*(X_{(i1)t+1} = x_{(i1)t+1} \mid X_{(i1)t} = x_{(i1)t}, n_{it}) = \\ &\sum_{x_{(i2)t}} F_i^\infty(x_{(i2)t} \mid x_{(i1)t}, n_{it}; \{CCP\}^*) F_i(x_{(i1)t+1} \mid x_{(i1)t}, x_{(i2)t}, n_{it}; \{CCP\}^*). \end{aligned} \tag{11}$$

To understand how equation (11) constrains beliefs, note first that the evolution of the state of plot $j \in \mathcal{M}_{(i1)}$ between t and $t+1$ depends on CCP_j^* at t. This investment decision rule depends, in turn, upon the state of j's adjacency at t, formed by plot j and all of its neighbors, including plot i. All of the plots in j's adjacency are in $\mathcal{M}_{(i1)}$ and $\mathcal{M}_{(i2)}$. Therefore, the state variables of plot owner j are contained in $\{n_{it}, X_{(i1)t}, X_{(i2)t}\}$.

[18]Weintraub et al. (2008), Benkard et al. (2015) and Ifrach and Weintraub (2017) consider alternative "oblivious equilibrium" concepts in the context of the Ericson and Pakes (1995) model of industry dynamics and show that they closely approximate the corresponding MPE. While we expect similar approximation results to hold in our setting, we leave this issue for future research.

[19]Sears et al. (2022b) apply a moment-based Markov equilibrium approach along the lines of Ifrach and Weintraub (2017) to estimate a model of groundwater extraction in California. Hodgson (2024), building on the work of Fershtman and Pakes (2012) on games of asymmetric information, also uses a reduced-state equilibrium concept to obtain a tractable empirical model of informational externalities in the context of strategic oil field exploration.

[20]Since the state of the map is an irreducible and aperiodic Markov chain, a unique stationary distribution exists.

If the owner of plot i knew $X_{(i2)t}$, they would thus be able to predict their neighbor j's behavior at t using CCP_j^* and, together with other primitives such as the drilling success and well failure processes, predict the stochastic evolution of the state of plot j which is the second factor in each term of the summation in equation (11). An AE, however, assumes that $X_{(i2)t}$ is not in plot owner i's information set but that they form expectations about it using the (conditional) steady state distribution $F^{\infty}(X_{(i2)t} \mid X_{(i1)t}, n_{it}; \{CCP\}^*)$ as probability weights on the RHS of equation (11).[21] Although each plot owner still has a unique CCP and set of beliefs, and the plot owners' joint decisions still depend on the entire cadastral map, the AE concept achieves considerable simplification.[22]

4 Structural Estimation

We now describe a tractable empirical structural model and a Simulated Method of Moments (SMM) procedure to estimate it. In addition to the four structural parameters $\Omega = (\theta, \alpha, \delta, \sigma)$, Table 1 summarizes the model primitives and how they are estimated.

4.1 Empirical specification

Unobserved heterogeneity We allow for two latent plot-specific factors. The first concerns the presence/absence of water-bearing fissures in the hard-rock underlying the plot, the distribution of which we approximate as Bernoulli. In other words, plots are either "developable" ($\mathcal{D} = 1$), in which case the model of optimal annual drilling choice outlined in Section 3 applies, or not ($\mathcal{D} = 0$), in which case our model is irrelevant for that plot and there is never drilling on it. To pin down the fraction of developable plots $\Pr(\mathcal{D} = 1)$, we exploit the fraction of "active" plots (see Section 2.2) observed in our sample as discussed below. The second latent factor applies only to developable plots and concerns heterogeneity in groundwater availability. Section 4.3 discusses estimation of the discrete distribution of this latent factor using well flow and failure data. We assume that farmers know the latent types of their own plot, but not those of the other plots in

[21] In our empirical implementation, we do not use equation (11) directly but rather compute equilibrium beliefs using "brute force" by simulating very long histories of investment, success, and failure events for every plot in each cadastral map until a steady state is reached. We then use simulated histories to compute the requisite transition probabilities.

[22] Using Brouwer's fixed point theorem, we can show that at least one AE exists. Multiplicity of equilibria, however, cannot be ruled out. Xu (2018) establishes that, in a static version of a similar model, the best response operator has a contraction property provided that the "strategic interaction parameter" is small enough. An extension of this result to a dynamic setting is nontrivial and is left as a topic for future research.

Table 1: ESTIMATION ROADMAP

	Symbol(s)	Section/Table note
Estimated in 2nd stage:		
Production function	θ, α, δ	4.4, 4.5
Scale of drilling shock	σ	4.4, 4.5
Fraction of developable plots	$\Pr(\mathcal{D} = 1)$	4.1, 4.4, 4.5
Estimated in 1st stage:		
Flow state probability functions	$\pi_1, ..., \pi_5$	4.3/note 1
Failure probability function	π_F	4.3/note 1
Flow/fail heterogeneity	high/low	4.3/note 2
Other primitives:		
Unsuccessful drilling cost	c_d	note 3
Successful drilling cost	c_s	note 4
Pr(Good monsoon)	$\mathbb{E}\, R_t$	note 5
$\Pr(\text{Success}\vert\mathcal{D} = 1)$	π_S	2.2/note 6
Discount factor	β	4.1
Plot network (by map-village)	$\{A, \mathbf{M}\}$	2.2/note 7

Notes: (1) See also Appendix F for estimation details. Probability functions depend on number of functioning borewells in adjacency and vary by mandal, monsoon rainfall and unobserved type; (2) Probability of low flow type = 0.346, high flow type = 0.654; (3) $c_d = 28,800$ Rs is median drilling cost (in 2017 Rs) across all borewells sunk since 2000; (4) $c_s = 59,800$ Rs is median of sum of drilling, pipe, casing, and electrical connection costs (in 2017 Rs) across all borewells sunk since 2000; (5) Good monsoon defined as June-November rainfall in mandal above study area average (See Appendix Figure F.1); (6) Observed drilling success per attempt = 0.375.; (7) At least one cadastral map per mandal (two mandals have two maps each; see Appendix B).

their adjacency. Furthermore, we assume that latent developable type and flow/failure type are independent of each other and that each is i.i.d. across plots.[23]

State space restrictions and observed types We assume that CCPs depend on reference plot area a and on the number, but not on the areas, of adjacent plots. This restriction effectively reduces $X_{(i1)t}$ to $\mathcal{N}_{it} = \sum_{j \in \mathcal{M}_{(i1)t}} n_{jt}$, yielding state-space $(\mathcal{N}_{it}, n_{it})$. Given the one-to-one mapping between plots, CCP *functions* and beliefs, and the very large number of plots in each village map, we do not allow for one CCP function specific to

[23]These i.i.d. assumptions rule out spatial clustering of drilling in areas more suitable for groundwater development. While positive spatial correlation in developable and flow/failure type is plausible, incorporating it is unlikely to appreciably improve the fit of the structural model. Given this, along with the limited information about spatial correlation in our data, as well as the additional complexity involved, we leave this refinement for future work.

every plot. Instead, for tractability we group similar plots into categories or types. Thus, there are (only) as many different CCPs as the number L of such types. Discretization of reference plot area into quintiles coupled with the number of adjacent plots in the maps ranging from 1 to 7, yields 35 possible observed types, which, along with 2 unobserved flow/fail types mentioned above, implies $L = 70$ for each of the 14 map villages, for a total of 980 distinct $CCP(\mathcal{N}_{it}, n_{it})$ functions in our empirical implementation.

Discount factor Given the challenge of identifying the discount factor in dynamic discrete choice models (Rust 1994; Magnac and Thesmar 2002), we follow standard practice and fix β in the SMM estimation. However, since our welfare calculations may be sensitive to this choice, instead of selecting a value of β from the literature, we calibrate it using data on land values. To do so, we first estimate the structural model at different (fixed) values of β along a coarse grid. Next, we simulate from each of these estimated models in steady state the average difference of the present discounted value per acre of active plots versus inactive plots. Lastly, we compare this value differential to its empirical counterpart, which we estimate to be around 80,000 Rs/acre (see Appendix D for details), and select the β that yields the closest match. This procedure delivers a value of $\beta = 0.95$.

4.2 Solution algorithm

Given values of Ω, $\Pr(\mathcal{D} = 1)$ and all the other primitives, we obtain an AE for each of the 14 cadastral maps (2,862 plots in total) as follows:

Initialize the maps:

Step 1 Draw a $\mathcal{D}_j$ for each plot j from the Bernoulli distribution with $\Pr(\mathcal{D} = 1)$.

Step 2 Assign each plot with $\mathcal{D}_j = 1$ an unobserved flow type ν_1 or ν_2, drawing from a Bernoulli distribution with probability of (low) type 1 = 0.346.[24]

Step 3 Assign each plot an initial number (zero) of functioning borewells $\{n_{j0} : j = 1, \ldots, P\}$ and an initial choice probability function (constant equal to 0.5) to each type $\{CCP_{l,0} : l = 1, \ldots, L\}$.

Iterate on beliefs and CCPs:

[24]To ensure a unique AE despite the inherent randomness of a particular map draw in Steps 1 and 2, we repeat these two steps ten times for each plot type and pool the resulting data in computing beliefs in Steps 4 and 5.

Step 4 Given $\{CCP_{l,q-1} : l = 1, \ldots, L\}$ at iteration $q = 1, 2, ...$, simulate the time-series of well drilling decisions, successes and (unobserved type-specific) failures in every plot on the map until the steady state is reached. Simulate $T = 150,000$ periods forward *in* steady state.

Step 5 From the steady state simulations, construct estimates of the one-period ahead state transition matrices $F(\mathcal{N}'|\mathcal{N}, n)$ for each type, averaging across plots on the map of the same type. Denote these estimates by $\widehat{F}_{lq}$.

Step 6 Given beliefs $\widehat{F}_{lq}$ and primitives, use value function iteration to compute new CCP's which solve the plot owner's game "against nature". Upon convergence of value function iteration, obtain a $\{CCP_{lq}\}$ satisfying the fixed point condition $V_{lq} = \Psi(V_{lq}, \widehat{F}_{lq}, \Omega)$ where V_l is the ex-ante integrated value function and Ψ is a value function iteration operator corresponding to the right-hand side of Bellman equations (4)-(10) for all types.

Convergence:

Step 7 If $\|CCP_q - CCP_{q-1}\|$ is small enough, then stop. If not, then update q and return to Step 4. If CCPs converge, so do beliefs, which are a continuous function of CCPs.

Steps 1-7 are nested within a routine for minimizing the SMM criterion function with respect to Ω and $\Pr(\mathcal{D} = 1)$ using a downhill simplex method.

4.3 First-stage: Estimating well interference effects

Appendix F lays out the panel data, econometric procedures, and results for our joint estimation of the well flow and failure processes incorporating interference from neighboring borewells. We use the 2017 adjacency survey to construct a 2012-16 panel of borewells at risk of failure and we combine the 2010 and 2017 household surveys into a two-year borewell flow panel. We allow for the endogeneity of neighboring borewells by letting their number be correlated with plot-specific unobserved heterogeneity in groundwater availability, our second latent factor. In other words, we exploit our panel data to estimate a correlated random effects model capturing how the number of functioning neighboring wells impact own well flow and failure *net* of common location-specific unobservables. In the case of failure, for example, our estimated interference effects can thus be interpreted as follows: Following a failure (or successful drilling attempt) *shock* to a neighbor's plot in period $t - 1$, there will be one less (more) surrounding well in period t, which lowers (raises) the probability that the own well fails in period t. This approach also allows us to recover predicted probability functions used in the second-stage estimation.

As mentioned in Section 2.2, because adjacency survey respondents may recall the functioning status of borewells belonging to their neighbors less accurately than of those on their own (reference) plots, our first-stage estimation procedure corrects misclassification error in the reported number of functioning wells on adjacent plots $\mathcal{N}_{it}$ in the failure component of the flow-failure model.[25] We use $\mathcal{N}_{it}^E$, the reported number of existing (functioning plus failed) wells, which is plausibly observed accurately, essentially as an instrument for the reported number of functioning wells. That is, we assume that the number of existing wells in the adjacency is associated with own-well failure only insofar as it proxies for the number of functioning wells in the adjacency, or, more formally, that own-well failure and $\mathcal{N}_{it}^E$ are orthogonal conditional on $\mathcal{N}_{it}$ and the random effect capturing the flow/fail type of the own-well.

In addition to conditioning the flow state and failure probabilities on the (endogenous) number of neighboring borewells, we also control for (exogenous) rainfall, mandal dummies, and plot area because well interference is less likely to affect larger plots ceteris paribus, as seen in Appendix Table F.4. We obtain adequate model fit with two unobserved plot types, low flow (high failure) with probability 0.346 and high flow (low failure) with probability 0.654. Figure 4 shows expected well flow $\sum_k \pi_k q_k$ (left panel), measured as a fraction of the outlet pipe's capacity and thus bounded above by 1, and the probability of well failure π_F (right panel) against number of functioning wells N, averaging across mandals and plot area quintiles for ease of presentation. While expected flow differs modestly between high and low unobserved flow types, the *marginal* effect of N on expected flow (the intensive margin externality) is virtually identical across types, whereas both the rate of well failure and the marginal effect of N on failure (extensive margin externality) are much higher for the low flow than high flow type. By contrast, an above average ("good") monsoon substantially increases well flow but has a statistically insignificant effect on failure.

[25] The flow component of the model uses a slightly different measure of $\mathcal{N}_{it}$, which is available only in the 2010 and 2017 household surveys. Because this measure is contemporaneous rather than retrospective, we assume no reporting error (see Appendix F for more detail).

Figure 4: Expected flow and failure probabilities

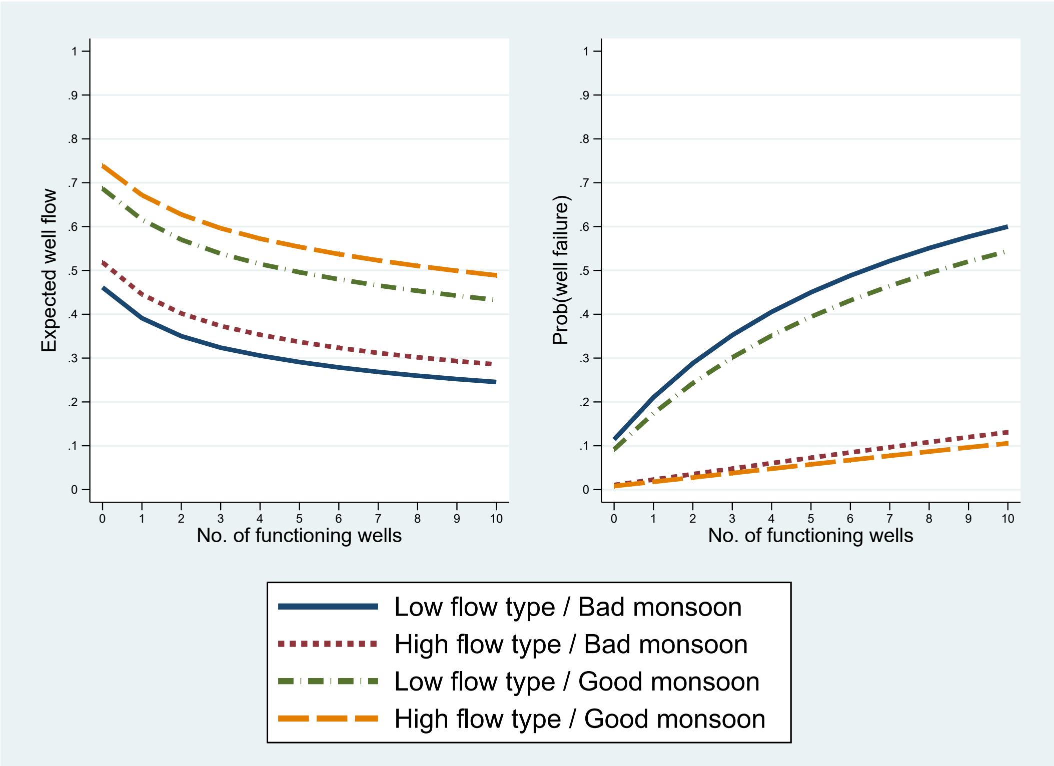

Notes: Left panel: Expected well flow ($\sum_k \pi_k q_k \leq 1$) as a function of N, the number of adjacent functioning wells, by latent flow type and monsoon state. Right panel: Annual probability of well failure as a function of N by latent flow type and monsoon state. Probability functions are predictions from the joint well flow/failure estimation averaged over the 12 mandals and 5 area quintiles.

4.4 Second-stage: Moment conditions and identification

We match the observed annual drilling rates in the representative sample of plots by area quintile a_k and number of functioning borewells to their model-based counterparts.[26] Since no investment occurs once a plot has two borewells by assumption, we do not match drilling rates conditional on $n \geq 2$. Although we do not target moments involving the number of functioning borewells outside of the reference plot (i.e., average drilling rates conditional on different values of $\mathcal{N}$), we later exploit the partial correlation between drilling and $\mathcal{N}$ to validate the model. This correlation is not needed in our second-stage because we estimate all parameters associated with well interference externalities in the first-stage, as previously discussed.

Identification of the structural parameters may be thought of, heuristically, in terms

[26]Model-based drilling rates are averages across the 14 map villages weighted by the proportion of total plot area in the representative plot sample contributed by sample plots associated with that village.

of a static model wherein drilling decisions are made once and for all, without borewell failure, and with the number of functioning borewells on a reference plot taken as given. In this case (ignoring unobserved heterogeneity),[27] we have $P_{n,k} \equiv \Pr(d = 1|n, a_k) = \text{logit}^{-1}(\{\theta[f(n+1, a_k; \alpha, \delta) - f(n, a_k; \alpha, \delta)] - \mathbb{E}\, c\}/\sigma)$, where $\mathbb{E}\, c = c_s\pi_s + c_d(1 - \pi_s)$ is the expected cost of drilling and $f = \frac{1}{\theta}\,\mathbb{E}[y(N, R, n, a_k; \alpha, \delta)]$ is expected (physical) output on a plot of area a_k with n functional borewells, where expectations are taken with respect to flow outcomes, beliefs about the number of functioning borewells in other adjacency plots, and rainfall.

Given α and δ, the difference in drilling rates across any two (n, k) pairs with different expected physical outputs identifies the ratio θ/σ.[28] Since $\mathbb{E}\, c$ is a known constant, we can then back out σ, and hence θ, from the average drilling rate at any (n, k). Given $\theta(\alpha, \delta)$ and $\sigma(\alpha, \delta)$, the remaining eight moment conditions yield more than enough equations to solve for α and δ. Intuitively, fixing a_k, differences in drilling rates at $n = 1$ and $n = 0$ capture diminishing returns to flow because, in log odds form, $\log \frac{P_{1,k}(1-P_{0,k})}{P_{0,k}(1-P_{1,k})} = \frac{\theta(\alpha,\delta)}{\sigma(\alpha,\delta)}\{[f(2, a_k; \alpha, \delta) - f(1, a_k; \alpha, \delta)] - [f(1, a_k; \alpha, \delta) - 0]\}$. Likewise, now fixing n, differences in drilling rates across area quintiles capture how the marginal product of flow varies with area because $\log \frac{P_{n,k}(1-P_{n,k'})}{P_{n,k'}(1-P_{n,k})} = \frac{\theta(\alpha,\delta)}{\sigma(\alpha,\delta)}\{[f(n+1, a_k; \alpha, \delta) - f(n, a_k; \alpha, \delta)] - [f(n+1, a_{k'}; \alpha, \delta) - f(n, a_{k'}; \alpha, \delta)]\}$.

Finally, to identify the fraction of latent developable plots $\Pr(\mathcal{D} = 1)$, we simulate the active status of a plot of each observed type using criteria analogous to those deployed in the actual data (Section 2.2); that is, we construct synthetic 5-year panels in steady state and assign an active status indicator $\mathcal{A}$ equal to one if any drilling attempt occurs over the panel or if there is a functioning borewell in the initial period. Averaging over plots of the same area type and over 5 simulated panels per type yields the model-based moments $Pr(\mathcal{A} = 1|a_k)$ that we match to those from the representative plot sample.

In all, we have 15 moment conditions and the SMM criterion function uses a diagonal weighting matrix consisting of the inverse of these moment variances. A small-scale Monte Carlo experiment reported in Appendix E indicates that our second-stage structural estimation algorithm performs exceedingly well in practice, recovering all of the true parameters with very small biases.

[27] While we incorporate unobserved (flow/fail) heterogeneity in the empirical model without the need to estimate the parameters of its distribution in the second-stage (see Table 1), abstracting from this heterogeneity delivers the simple intuitive mapping from drilling moments to structural parameters.

[28] For instance, using differences in log odds ratios, we obtain $\theta/\sigma = \log \frac{P_{0,k}(1-P_{0,k'})}{P_{0,k'}(1-P_{0,k})}[f(1, a_k; \alpha, \delta) - f(1, a_{k'}; \alpha, \delta)]^{-1}$.

4.5 Results and model validation

Table 2 reports the second-stage estimates along with their standard errors based on a 100 replication bootstrap.[29] We strongly reject a Cobb-Douglas production function (i.e., $\delta = 0$) in favor of a CES. Figures 5 and 6 show that our model matches the 15 targeted moments reasonably well. As a goodness-of-fit test, consider the functioning borewell density implied by the estimated model in steady state of 0.229 wells per developable acre or 0.153 ($= 0.229 \times 0.67$) wells per acre of land overall. This latter figure is quite close to the 2012-16 average of 0.160 borewells per acre in our representative plot sample.[30]

Table 2: Structural Parameter Estimates

θ	α	δ	σ	$\Pr(\mathcal{D}=1)$
16.99	0.78	0.72	0.63	0.67
(0.28)	(0.01)	(0.04)	(0.03)	(0.02)

Notes: Bootstrapped standard errors in parentheses. See equation (1) for definition of production function parameters (θ, α, δ); σ is scale of drilling shock.

By way of validation, we simulate data from the estimated model to compute the partial correlation between d_{it} and $\mathcal{N}_{it}$ and compare it to its data counterpart. This "reduced-form" parameter captures the strength of strategic substitutability between neighbors' drilling decisions. Starting at a steady state on each of 10 replications of the cadastral village maps, we simulate five-year panels consisting of triplets $\{d_{it}, n_{it}, \mathcal{N}_{it} : t = 1, ..., 5\}$ for every plot on the map that is assigned developed status (see Step 2 in Subsection 5.3); this yields 61,695 5-year panels in total. On this large "sample", we estimate a linear probability drilling model with reference plot fixed effects using only observations with $n_{it} < 2$, i.e., the cases when drilling is theoretically possible. We run two different versions of this regression, one conditioning on n_{it} and the other not. Since we can make the sample arbitrary large, each regression yields a "population" value of the strategic substitutability effect. In Table 3, we compare each model-derived population value with its empirical counterpart estimated from the adjacency panel data (see Appendix Table H.1) using a bootstrap hypothesis test suggested by Hansen (2022). Despite reasonably tight percentile-t confidence intervals, for neither specification of strategic substitutability can we reject the null hypothesis of equality between model and data.

[29]These standard errors are understated as they do not account for pre-estimated first-stage parameters.

[30]Following a referee's suggestion , we also assess whether incorporating additional heterogeneity in land productivity, specifically in θ, would better fit the observed correlation between plot values and the number of functioning borewells. Our conclusion, reported in Appendix G, is that it would not; the heterogeneity built into our model (more than) fully accounts for this external data moment.

Figure 5: Annual drilling rates by plot area quintile and n

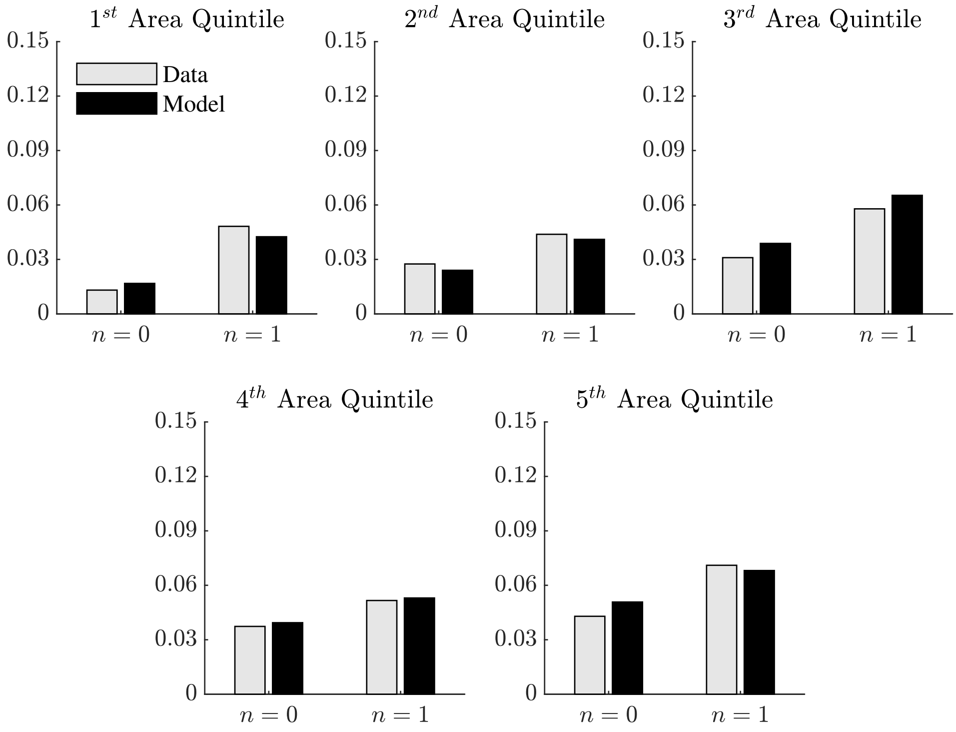

Notes: Each pair of bars represents a data moment (annual drilling rate by plot area quintile and number of functioning wells on the reference plot) and its corresponding model fit.

Figure 6: Probability plot is active by area quintile

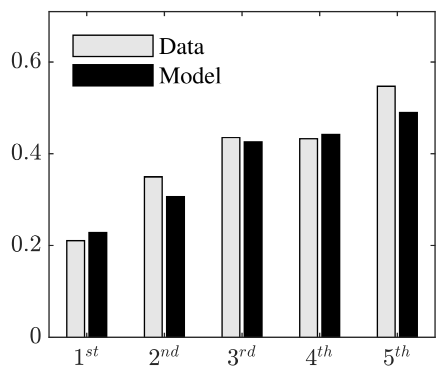

Notes: Each pair of bars represents a data moment (probability of having at least one functioning borewell or drilling attempt in a 5-year period by plot area quintile) and its corresponding model fit.

Table 3: STRATEGIC INTERACTIONS – DATA VS. MODEL

Conditioning on n	Data	Model	H_0 : Data = Model bootstrap p-value
No	-0.0442 [-0.007,-0.0882]	-0.0363	0.691
Yes	-0.0381 [-0.0008,-0.0801]	-0.0340	0.845

Notes: Percentile-t bootstrap confidence intervals in square brackets. Row 1 specification regresses d_{it} on $\mathcal{N}_{it}$ (cf. column 5 of Appendix Table H.1); row 2 specification regresses d_{it} on $\mathcal{N}_{it}$ and a dummy for $n_{it} = 1$. (cf. column 6 of Appendix Table H.1). All regressions include plot fixed effects.

5 Counterfactuals

Our quantitative policy evaluation addresses four questions: (1) What is the social cost of the current policy of free (but rationed) electricity to farmers for pumping groundwater? (2) What is the optimal tax on borewells that eliminates the deadweight loss of electricity subsidies and of the well interference externality? (3) How would the burden of borewell taxation be distributed across landowners? (4) How much do local externalities magnify the inefficiency of electricity subsidies and distort the spatial allocation of drilling?

5.1 Social cost of electricity subsidies and optimal tax

We estimate the private value of developable land, or private welfare for short, to be 54,700 Rs/acre.[31,32] Social welfare is private welfare net of the cost of electricity, which, though given free to borewell owners, is not free to society. The steady-state fiscal cost of the electricity subsidy in present value terms is 38,900 Rs/acre = 0.229 borewells per acre $\times$ 8,500/0.05, where the second term is the annual cost of electricity to run a pump given daily power rationing (Appendix A) divided by the annual discount rate. Thus, social welfare is only 15,800 Rs/acre; more than two-thirds of private welfare is accounted for by the capitalized value of the electricity subsidy.

Consider next an annual tax τ on functioning borewells, which could be implemented in practice as a flat charge for maintaining an agricultural electrical connection. Setting τ equal to the annual cost of electricity (i.e., τ_e =8,500 Rs/year) would fully recover costs from agricultural consumers. A flat charge exceeding τ_e would act, at the margin, like a Pigouvian tax on borewells.[33]

[31] This is the average expected discounted present value of agricultural output minus drilling costs in steady state across all maps weighted by the proportion of total acreage from the map village represented in the sample (note that we do not include the extreme-value random shocks in our valuations). Undevelopable land has zero private value by assumption (our normalization). To be clear, private welfare accounts for externalities inasmuch as it averages values across hundreds of adjacent plots with (potentially) mutually interfering borewells.

[32] The value of developable land (i.e., over and above the value of undevelopable land) is much lower than the value of *active* land, which we estimated at 80,000 Rs/acre, because the proportion of developable plots (0.67) is much higher than that of active plots (0.39).

[33] In terms of the model, the annual *net* value of output in Rs under a counterfactual tax $\tau > 0$ becomes $\mathbb{E}\, y - \tau n$. Once a borewell fails, its owner incurs no further tax on it. We further assume that extant borewells that become unprofitable under a new tax are dismantled, i.e., their pipe and casing removed, at zero cost. Although our structural model does not explicitly incorporate dismantling (which would never be chosen anyway), our counterfactuals still treat this decision as the outcome of a strategic equilibrium (Appendix I). In particular, we assume that once a tax is implemented, borewell owners adopt beliefs (one-period ahead state transition probabilities) consistent with the transition to the new steady state under the counterfactual policy (see below). Thus, in making their dismantling decision, each borewell owner takes into account the dismantling done by other borewell owners in their own adjacency.

For each counterfactual τ, we compute social welfare along the entire transition path from the zero-tax baseline to the new steady state AE. A key assumption is that each plot owner at each date of the transition has correct beliefs about the evolution of the state of their adjacency when averaged over their respective village map's stochastic transition paths (see Appendix I for details of the equilibrium concept and solution algorithm). Although this rationality assumption is stronger than the one invoked to compute the steady state AE, calculating social welfare along the transition path is preferable to simply comparing long-run steady states as it accounts for the short-run benefits and costs that accrue to *current* landowners as they transition from their initial stock of wells without a tax to the steady state with a tax. As Domeij and Heathcote (2004) argues in a related context, ignoring the transition path in welfare calculations can be misleading for policy.

Figure 7: Welfare, borewell density, and output under alternative taxes

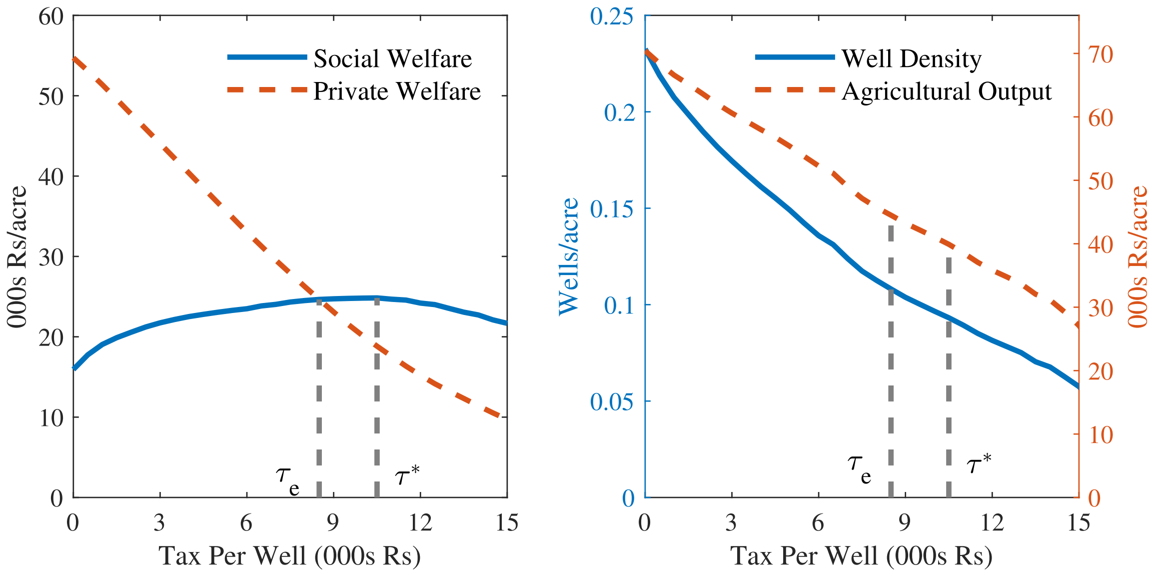

Notes: In the left panel, each point on the solid (dashed) curve represents the social (private) per acre value of developable land, or welfare, along the transition path from the benchmark zero-tax economy to the long-run steady-state under alternative values of τ: $\tau_e = 8.5$ is the tax that recovers electricity costs; $\tau^* = 10.5$ is the optimal tax. In the right panel, each point on the curve represents, respectively, borewells per acre and the present discounted value of agricultural output in the long-run steady-state under alternative values of τ.

Figure 7 plots social and private welfare (left panel) and borewell density and the present discounted value of agricultural output per acre (right panel) against the annual borewell tax. Private welfare, borewell density, and output all decrease monotonically as τ increases, with output declining at a slower rate than density as drilling increasingly concentrates on high flow/low failure plots. By contrast, social welfare is hump-shaped in τ and intersects private welfare exactly at τ_e, the social cost of electricity.

Deadweight loss of electricity subsidies Arguing that usage-based pricing of agricultural power is impractical in India, Shah et al. (2007) propose instead a flat-fee combined with quantity rationing (the latter discussed by Ryan and Sudarshan 2022). Implementing this policy in the form of an annual borewell tax $\tau = \tau_e$ increases the social value of groundwater development from 15,800 to 24,600 Rs/acre.[34] This implies a deadweight loss from free electricity provision of 8,800 Rs (around 135 US$) per developable acre in present value terms, or 23% of the fiscal cost of the subsidy. In other words, nearly a quarter of every rupee transferred in-kind to farmers in the form of electricity is lost through over-drilling.

Optimal borewell taxation To correct the negative externality, the social welfare maximizing tax τ^* should exceed the annual cost of electricity τ_e.[35] In the left panel of Figure 7, we find that social welfare (solid curve) is maximized at $\tau^* = 10,500$ Rs, which is about 24% higher than τ_e. Nevertheless, there is little welfare difference between a tax of 8,500 Rs and one of 10,500 Rs, and only a small absolute difference in equilibrium borewell density (right panel). Once electricity is provided to farmers at cost, and borewell density declines by 52% from its baseline value as a result, the *marginal* externality cost is small.

5.2 Distributional implications of borewell taxation

Accounting for transitional dynamics, as we do in this paper, also allows us to understand how the burden of borewell taxation is distributed across landowners, which can inform the policy debate. At each counterfactual tax τ, Figure 8 shows private welfare by the number of functioning borewells on the plot (n) *before* the introduction of the policy (private welfare in the left panel of Figure 7 is a weighted average of these three conditional curves). Note that taxing borewells, regardless of τ, has practically zero impact on plots with no functioning borewells today because the option value of future drilling is minimal; these developable plots are strongly selected to be of the low flow/high failure type.

[34]Incremental output would also decline by 37% in present value terms. While dry season production contributes only a fraction of annual output (about a third in the case of paddy, the principal crop), prices could potentially increase was cost-recovery implemented on a district or state-wide scale. India's high degree of spatial agricultural market integration (e.g., Ghosh, 2011), however, suggests that it is reasonable to assume fixed crop prices in our counterfactuals.

[35]Optimality of a centralized Pigouvian tax presupposes that landowners cannot restrain socially undesirable drilling through side-payments to neighbors. Aside from enforcement issues, one argument against this Coaseian solution in our setting is the complex, multi-lateral, negotiations it would require across the entire network of agricultural plots.

Figure 8: Private welfare by initial number of functioning borewells

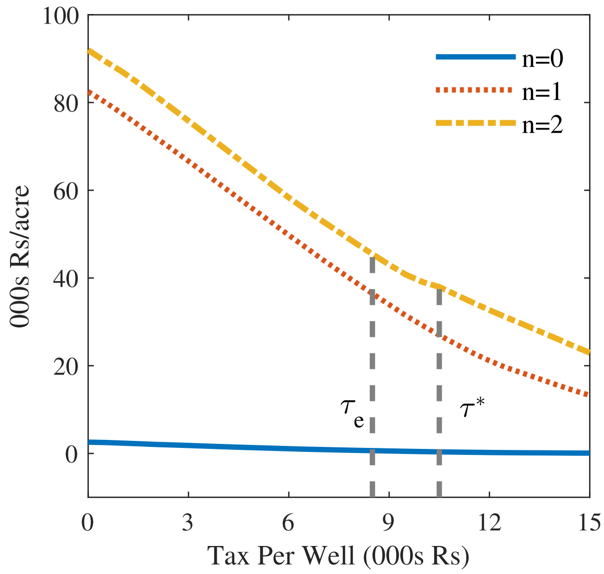

Notes: Each point represents private welfare along the transition path from the benchmark zero-tax economy to the long-run steady-state under alternative values of τ for different values of n at the time the policy is introduced: $\tau_e = 8.5$ is the tax that recovers electricity costs; $\tau^* = 10.5$ is the optimal tax.

Turn now to the cost-recovery tax τ_e. Table 4 displays private welfare with 0, 1, and 2 functioning borewells at baseline and, respectively, under this counterfactual. As noted above, the welfare loss for $n = 0$ plots is negligible, whereas owners of $n = 1$ plots lose 147,000 Rs and owners of $n = 2$ plots lose 237,000 Rs. There is considerable plot type heterogeneity *within* each of the $n = 0, 1, 2$ groups, which is costly (if not impossible, in the case of the two latent factors) for the policy-maker to observe. By contrast, n is observable because a functioning borewell requires a working electrical connection. Our analysis thus suggests the contours of a practical compensation scheme to mitigate political opposition to electricity cost recovery. Farmers would receive 130,000 Rs per existing agricultural connection upfront if they agree to pay τ_e annually as long as their connection is active; those who successfully drill and establish a connection ex-post would not be compensated. The last column of Table 4 shows relatively small welfare losses or even gains under this compensation formula. Furthermore, the ratio of compensation to fiscal cost of the subsidy is only 0.765 ($= 130/\ (8.5/.05)$), which means that the government would save about 23% on its budget for electricity provision while leaving farmers largely whole (reflecting our earlier finding that the deadweight loss from free electricity is 23% of the fiscal cost). A compensation scheme based on n alone is surprisingly efficacious.

Table 4: Compensation Scheme for Electricity Cost Recovery

	Private welfare				
Plot type	Baseline	Cost-recovery	Δ	Compensation	Net welfare
$n = 0$	5.7	1.4	-4.3	0	-4.3
$n = 1$	263.1	116.3	-146.8	130.0	-16.8
$n = 2$	468.1	231.4	-236.7	260.0	23.3

Notes: All figures are in thousands of Rs per plot.

5.3 Local externalities, deadweight loss and over-drilling

To isolate the role of local externalities, we compute the village map equilibria in a counterfactual "island" economy with zero well interference effects. For the sake of comparability, we reduce the TFP parameter θ in the island economy to deliver the same overall borewell density (0.229 wells/acre) and, consequently, an identical fiscal cost of the electricity subsidy as in the baseline "network" economy. Because the profitability of drilling is substantially reduced in the island economy to equalize borewell density, social welfare falls to just 9,500 Rs/acre (compare columns 1 and 3 of Table 5). Removing electricity subsidies in this counterfactual economy, however, increases social welfare by 4,700 Rs per acre as compared to 8,800 Rs per acre in the network economy. We conclude that well interference nearly doubles the deadweight loss arising from provision of free electricity to borewell owners.

Table 5: Electricity Cost Recovery Counterfactuals

	Network economy		*Island economy*	
	Baseline	Cost-recovery	Baseline	Cost-recovery
Borewells density (wells/acre)	0.23	0.11	0.23	0.07
Fiscal cost of subsidy ('000 Rs/acre)	38.9	0.0	38.9	0.0
Social welfare ('000 Rs/acre)	15.8	24.6	9.5	14.2
Deadweight loss ('000 Rs/acre)	8.8		4.7	
DWL/Fiscal cost of subsidy	0.23		0.12	

Notes: For the island economy, we set $\theta = 11.15$, versus 16.99 in the network economy, so as to obtain equivalent borewell densities.

The network/island economy comparison also illustrates the influence of local externalities on where drilling occurs.[36] We simulate steady-states with and without externalities on a single map-village (Ayyavaripalli; see Appendix J) so as to focus on plot-level het-

[36] We thank an anonymous referee for suggesting this line of inquiry.

erogeneity, specifically, area and groundwater quality (flow/fail type). We again reduce θ in the island economy so as to equate its borewell density to that of the network economy (for this single map-village). At an average annual rate of 0.028, drilling is dramatically lower in the island economy as compared to the rate of 0.127 in the network economy. Figure 9 shows predicted plot-level differences in, respectively, number of borewells and drilling rates between network and island economies by area quintile and flow/fail type. Externalities 'reallocate' drilling activity and, consequently, functioning borewells from large plots of low flow/high failure type to large plots of high flow/low failure type. This reallocation is so pronounced that, despite the much higher drilling rate in the network economy, there is less drilling and fewer borewells *overall* on large plots of low flow/high failure type with externalities than without.

Figure 9: Allocation of Over-drilling

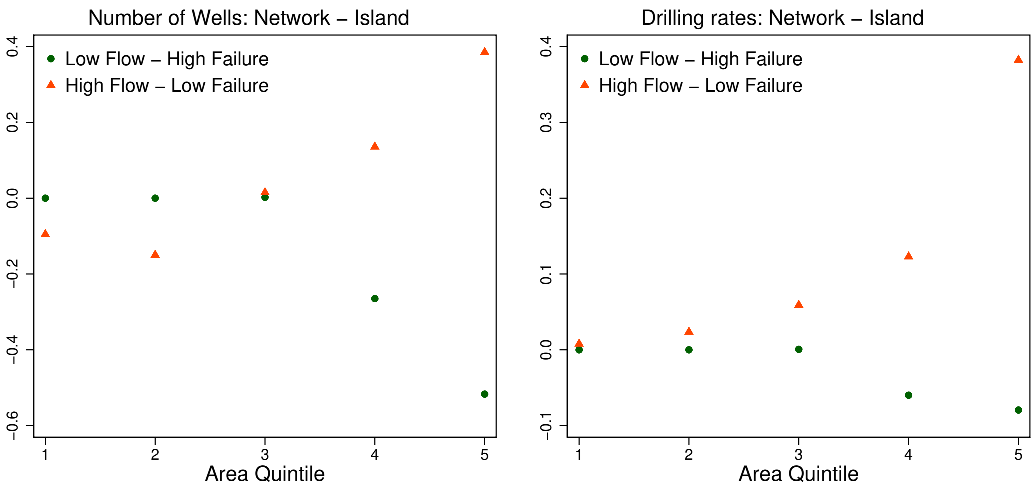

Notes: Predictions from a plot-level regression of differences in number of wells (drilling rate) between network and island economies on plot area quintile dummies, flow/fail type dummy, and their interactions. Since the simulated data set can be made arbitrary large, no confidence intervals are reported.

6 Conclusion

We set out to assess the social cost of a policy of free electricity to farmers for groundwater pumping in South India, a context with economically important yet highly localized externalities. To do so, we developed a tractable dynamic strategic equilibrium model of borewell investment across a large network of heterogeneous agricultural plots along with a novel simulation-based estimation strategy, one potentially applicable to a wide range of settings beyond that of concern in this paper.

Despite daily power rationing that obviates over-pumping (as per Ryan and Sudarshan 2022), once the extensive margin and drilling costs are taken into account, we find that subsidizing electricity is a rather inefficient means of transferring resources to farmers; almost a quarter of the transfer value is wasted through over-drilling of borewells, a distortion greatly exacerbated by well interference. While the optimal Pigouvian tax on borewells exceeds the annual cost of electricity by 24%, the marginal social gain from such a tax beyond that required for cost-recovery is minimal, reflecting the extensive dismantling of unproductive borewells and choking-off of drilling that cost-recovery would entail. Finally, although charging fully for electricity would have substantial short-run distributional implications, these could largely be ameliorated through a (practical) compensation scheme.

A limitation of our methodology is that it relies on a steady state assumption; it may, therefore, not carry over to settings in which water tables are undergoing significant decline, as in much of northwest India. Adapting our approach to these circumstances would be a fruitful area of future research.

References

Acemoglu, Daron, Camilo Garcia-Jimeno, and James A Robinson, "State capacity and economic development: A network approach," *American Economic Review*, 2015, *105* (8), 2364–2409.

Aguirregabiria, V. and P. Mira, "Dynamic discrete choice structural models: A survey.," *Journal of Econometrics*, 2010, *156* (1), 38–67.

Badiani-Magnusson, Reena and Katrina Jessoe, "Electricity Prices, Groundwater, and Agriculture: The Environmental and Agricultural Impacts of Electricity Subsidies in India," in "Agricultural Productivity and Producer Behavior," University of Chicago Press, 2018, pp. 157–183.

Benkard, C. Lanier, Przemyslaw Jeziorski, and Gabriel Y. Weintraub, "Oblivious equilibrium for concentrated industries," *The RAND Journal of Economics*, 2015, *46* (4), 671–708.

Blakeslee, David, Ram Fishman, and Veena Srinivasan, "Way down in the hole: Adaptation to long-term water loss in rural India," *American Economic Review*, 2020, *110* (1), 200–224.

Bruno, Ellen M and Katrina Jessoe, "Missing markets: Evidence on agricultural groundwater demand from volumetric pricing," *Journal of Public Economics*, 2021, *196*, 104374.

Burlig, Fiona, Louis Preonas, and Matt Woerman, "Energy, groundwater, and crop choice," *NBER Working Paper*, 2021, (w28706).

Carleton, Tamma, Levi Crews, and Ishan Nath, "Agriculture, trade, and the spatial efficiency of global water use," Technical Report 2023.

Chandrakanth, M.G., *Externality in Irrigation*, New Delhi: Springer India,

Cole, Shawn, Xavier Giné, Jeremy Tobacman, Petia Topalova, Robert Townsend, and James Vickery, "Barriers to household risk management: Evidence from India," *American Economic Journal: Applied Economics*, 2013, *5* (1), 104–35.

Domeij, David and Jonathan Heathcote, "On the distributional effects of reducing capital taxes," *International economic review*, 2004, *45* (2), 523–554.

Donna, Javier D and José-Antonio Espín-Sánchez, "The illiquidity of water markets," *Review of Economic Studies*, 2023.

Ericson, Richard and Ariel Pakes, "Markov-perfect industry dynamics: A framework for empirical work," *The Review of economic studies*, 1995, *62* (1), 53–82.

Fershtman, Chaim and Ariel Pakes, "Dynamic games with asymmetric information: A framework for empirical work," *The Quarterly Journal of Economics*, 2012, *127* (4), 1611–1661.

Fishman, Ram Mukul, Tobias Siegfried, Pradeep Raj, Vijay Modi, and Upmanu Lall, "Over-extraction from shallow bedrock versus deep alluvial aquifers: Reliability versus sustainability considerations for India's groundwater irrigation," *Water Resources Research*, 2011, *47* (6).

Fishman, Ram, Xavier Giné, and Hanan G Jacoby, "Efficient irrigation and water conservation: Evidence from South India," *Journal of Development Economics*, 2023, p. 103051.

Ghosh, Madhusudan, "Agricultural policy reforms and spatial integration of food grain markets in India," *Journal of Economic Development*, 2011, *36* (2), 15.

Giné, Xavier and Hanan G Jacoby, "Contracting under uncertainty: Groundwater in South India," *Quantitative Economics*, 2020, *11* (1), 399–435.

Hagerty, Nick, "What holds back water markets? Transaction costs and the gains from trade," *Unpublished, Working Paper*, 2023.

Hansen, Bruce, *Econometrics*, Princeton University Press, 2022.

Hodgson, Charles, "Information externalities, free riding, and optimal exploration in the uk oil industry," Technical Report, National Bureau of Economic Research 2024.

Hora, Tejasvi, Veena Srinivasan, and Nandita B Basu, "The groundwater recovery paradox in South India," *Geophysical Research Letters*, 2019, *46* (16), 9602–9611.

Ifrach, Bar and Gabriel Weintraub, "A framework for dynamic oligopoly in concentrated industries," *Review of Economic Studies*, 2017, (6), 1106–1150.

Jacoby, H. G., ""Well-fare" Economics of Groundwater in South Asia.," *World Bank Research Observer*, 2017, *32* (1), 1–20.

Jain, Meha, Ram Fishman, Pinki Mondal, Gillian L Galford, Nishan Bhattarai, Shahid Naeem, Upmanu Lall, Balwinder-Singh, and Ruth S DeFries, "Groundwater depletion will reduce cropping intensity in India," *Science advances*, 2021, *7* (9), eabd2849.

König, Michael D, Dominic Rohner, Mathias Thoenig, and Fabrizio Zilibotti, "Networks in conflict: Theory and evidence from the great war of africa," *Econometrica*, 2017, *85* (4), 1093–1132.

Kumar, M Dinesh, MVK Sivamohan, V Niranjan, and Nitin Bassi, "Groundwater management in Andhra Pradesh: time to address real issues," *Occasional paper*, 2011, *4*.

Lin, C. Y. Cynthia, "Strategic Decision-Making with Information and Extraction Externalities: A Structural Model of the Multi-Stage Investment Timing Game in Offshore Petroleum Production.," *Review of Economics and Statistics*, 2013, *95* (5), 1601–1621.

Magnac, Thierry and David Thesmar, "Identifying dynamic discrete decision processes," *Econometrica*, 2002, *70* (2), 801–816.

Pfeiffer, L. and C. Y. C. Lin, "Groundwater pumping and spatial externalities in agriculture.," *Journal of Environmental Economics and Management*, 2012, *64* (1), 16–30.

Rafey, Will, "Droughts, deluges, and (river) diversions: Valuing market-based water reallocation," *American Economic Review*, 2023, *113* (2), 430–471.

Rust, John, "Structural estimation of Markov decision processes," *Handbook of econometrics*, 1994, *4*, 3081–3143.

Ryan, Nicholas and Anant Sudarshan, "Rationing the commons," *Journal of Political Economy*, 2022, *130* (1), 210–257.

Sayre, Susan Stratton and Vis Taraz, "Groundwater depletion in India: Social losses from costly well deepening," *Journal of Environmental Economics and Management*, 2019, *93*, 85–100.

Sears, Louis S and C-Y Cynthia Lin Lawell, "Water management and economics," in "The Routledge Handbook of Agricultural Economics," Routledge, 2018, pp. 269–284.

— , — , **and M Todd Walter**, "Groundwater Under Open Access: A Structural Model of the Dynamic Common Pool Extraction Game," *unpublished manuscript, Cornell University*, 2022.

— , **C-Y Lin Lawell, Gerald Torres, and M Todd Walter**, "Moment-based Markov equilibrium estimation of high-dimension dynamic games: An application to groundwater management in California," 2022.

Sekhri, Sheetal, "Wells, water, and welfare: the impact of access to groundwater on rural poverty and conflict," *American Economic Journal: Applied Economics*, 2014, *6* (3), 76–102.

— , "Agricultural trade and depletion of groundwater," *Journal of Development Economics*, 2022, *156*, 102800.

Shah, Tushaar, *Taming the anarchy: Groundwater governance in South Asia*, Routledge, 2010.

— , **Christopher Scott, Avinash Kishore, and Abhishek Sharma**, "11 Energy–Irrigation Nexus in South Asia: Improving Groundwater Conservation and Power Sector Viability," *THE AGRICULTURAL GROUNDWATER REVOLUTION*, 2007, p. 211.

— , **Mark Giordano, and Aditi Mukherji**, "Political economy of the energy-groundwater nexus in India: exploring issues and assessing policy options," *Hydrogeology Journal*, 2012, *20* (5), 995–1006.

Sidhu, Balsher Singh, Milind Kandlikar, and Navin Ramankutty, "Power tariffs for groundwater irrigation in India: A comparative analysis of the environmental, equity, and economic tradeoffs," *World Development*, 2020, *128*, 104836.

Weintraub, Gabriel Y, C Lanier Benkard, and Benjamin Van Roy, "Markov perfect industry dynamics with many firms," *Econometrica*, 2008, *76* (6), 1375–1411.

Xu, Haiqing, "Social interactions in large networks: A game theoretic approach," *International Economic Review*, 2018, *59* (1), 257–284.

Online Appendix

Contents

A Average annual electricity costs per borewell

The electricity cost of a borewell per year is the product of (1) power consumption of the average pump of 6 horsepower (HP), which is 4.5 kWh (= 6 HP×0.746 kWh/HP), (2) 630 annual hours of pumping (average of unit record data for our 12 sample mandals from India's 4th Minor Irrigation Census), and (3) marginal cost of electricity of 3 Rs/kWh (Gulati and Pahuja, 2012). All three components in this calculation are likely overly conservative estimates, so that 8500 Rs. should be viewed as a lower bound on the true electricity cost.

B Cadastral map villages

The villages for which we have cadastral maps are Pamireddypalli in Atmakur mandal, Dharmapur and Ramachandrapuram in Mahabubnagar mandal, Jajapur in Narayanapet and Thipparasipalli in Utkur madal. In Anatapur district, we have cadastral maps for Manesamudram in Hindupur mandal, M. Venkata Puram and Manepalli, both part of the same panchayat in Lepakshi mandal, Y.B. Halli in Madakasira Muddireddy Palli in Parigi Chalakuru and Somandepalli, both part of the same panchayat in Somandepalli mandal, Siddarampuram and Reddipalli in B.K. Samudram mandal, Itukalapalli in Anantapur and Ayyavaripalli in Rapthadu mandal.

We take these maps to be representative of all villages for which we have data in each respective mandal. We use the digitized maps, as in the example shown in Figure B.1, to create 14 plot adjacency matrices defining the networks upon which the dynamic discrete investment game is played.

Figure B.1: Village Muddiredipalle

(a) Original Cadastral Map

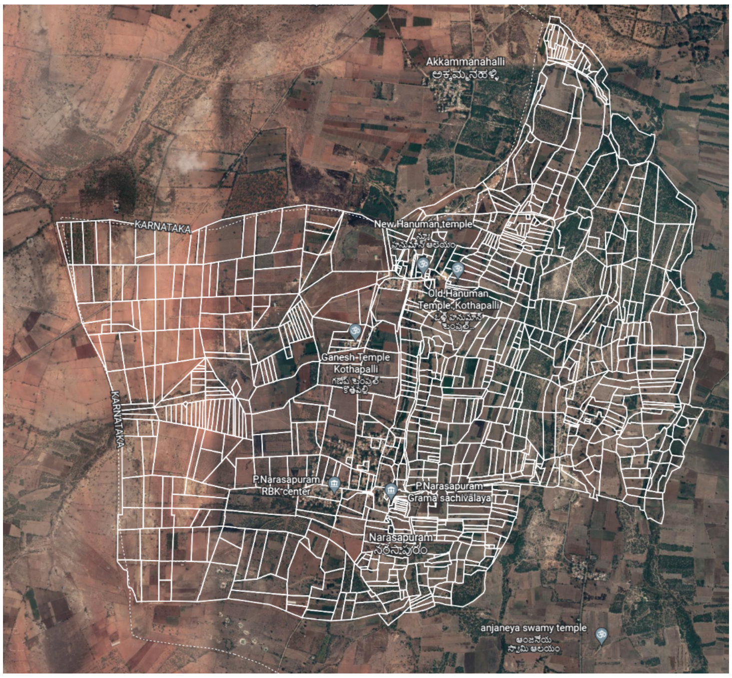

(b) Digitized Cadastral Map

C Wealth, liquidity and drilling

Using the 2010 household survey, we compute the total value of household assets in 2009, including agricultural land, livestock, agricultural machinery, household durable goods, and savings in the form of bank deposits, cash and jewelry. Gross wealth should be a good proxy for liquidity and thus allows us to test for the importance of financial constraints on well-drilling. Table C.1 presents estimates of the determinants of drilling attempts using the 2012-16 representative plot panel. Columns 1-3 report linear probability models with standard errors clustered at the plot level whereas columns 4-6 report plot-level random effects logit models. Inferences about the association between pre-sample wealth and drilling over the 5-year panel from the two sets of estimates are virtually identical.

In column 1 (and 4), which includes only mandal dummies as additional controls, we see that wealthier households (in 2009) are significantly more likely to drill during 2012-2016. However, once we also control for plot size with area quintile dummies (cols. 2 and 5), this effect largely disappears. Evidently, wealth and plot area are positively correlated and there is a greater incentive to drill on larger plots. Finally, in column 3 (and 6), we control for the initial number of functioning borewells on the plot entering the panel period, which further attenuates the wealth effect toward zero. Admittedly, these liquidity effects, or lack thereof, are associations, i.e., not necessarily causal. For example, despite controlling for initial borewells, unobserved (to us) suitability for groundwater development (and thus drilling) may have augmented pre-sample wealth through past groundwater use. Such correlations, however, would bias liquidity effects away from zero, contrary to what we find in Table C.1.

Further, according to our 2017 survey of borewell owners, only 27% of respondents relied on their own savings as the main source of finance for the largest components of the cost of borewell investment, drilling the bore and purchase of the submersible pump. Most farmers use various forms of formal and informal credit. More broadly, data from the household credit module contained in the 2010 household survey show that, out of the 1488 respondents, 89% had outstanding bank credit, 46% had loans from family or friends, and 45% were borrowing from moneylenders. These percentages are very similar across households that have at least one borewell on their land and those that have none. This substantial credit access in our setting may explain why pre-sample wealth is uncorrelated with subsequent drilling and provides empirical justification for a borewell investment model that abstracts from financial frictions.

Table C.1: Determinants of drilling

	LPM			RE Logit		
	(1)	(2)	(3)	(4)	(5)	(6)
Log(gross wealth 2009)	0.00531***	0.00236	0.00130	0.161***	0.0701	0.0379
	(0.00189)	(0.00193)	(0.00191)	(0.0578)	(0.0622)	(0.0622)
Plot area quintile (1st omitted):						
2nd	—	0.0110**	0.00960**	—	0.563***	0.519***
		(0.00442)	(0.00444)		(0.197)	(0.197)
3rd	—	0.0210***	0.0185***	—	0.847***	0.768***
		(0.00556)	(0.00551)		(0.203)	(0.203)
4th	—	0.0222***	0.0199***	—	0.877***	0.808***
		(0.00518)	(0.00520)		(0.195)	(0.194)
5th	—	0.0284***	0.0234***	—	1.024***	0.904***
		(0.00645)	(0.00655)		(0.213)	(0.213)
Initial borewells (0 omitted):						
1	—	—	0.0197***	—	—	0.549***
			(0.00471)			(0.117)
2	—	—	0.0433***	—	—	0.859***
			(0.0125)			(0.189)
Observations	14,310	14,310	14,310	14,310	14,310	14,310
Plots	2,862	2,862	2,862	2,862	2,862	2,862

Notes: Standard errors in parentheses (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$), clustered at plot-level (columns 1-3). Column 1-3 use a linear probability model; columns 4-6 use plot-level random effects logit MLE. Dependent variable is whether a drilling attempt was made on plot during the year (from 2012-16). Each regression also controls for mandal dummies (coefficients not reported).

D Land values and active status

We use plot value data collected in the 2017 household survey to estimate the difference in present discounted values between active and inactive land. Recall that an active plot is one on which at least one drilling attempt was made during 2012-16 or which already had at least one functioning borewell in 2012. Our survey asked each plot owner "if you were to sell this plot today, including the associated water rights, how much would you receive in thousands of Rs per acre?" In evaluating the present discounted value of the projected income flows off of their land, as reflected in their stated sales price, we presume that farmers use the same discount factor that they would use in assessing the future net benefits from drilling. This presumption is the basis for our calibration of β for the structural estimation.

To estimate the average marginal value of an active plot, we regress the reported value per acre of plot j (winsorized at 5% level) on the active status indicator $\mathcal{A}_j$. Needless to say, $\mathcal{A}_j$ is potentially endogenous. For instance, unobserved land amenities (e.g., ready access to markets, good soil) may both increase land values and encourage groundwater development. It is also plausible that poorer households both own less valuable land and can less afford to develop their land for groundwater extraction. To deal with such reverse causality, we focus on households owning multiple plots in the same village and use household fixed effects to estimate the land value regression. This procedure controls for both unobservable household-specific and location specific factors.

In Table D.1, we report three sets of two regressions, one without and one with controls for plot area quintile. The first set of regressions is by OLS, the second with survey village fixed effects (44 sample villages), and the third with household fixed effects (895 multi-plot households; 502 single-plot households).

Table D.1: Plot values and active status

	(1)	(2)	(3)	(4)	(5)	(6)
Active ($\mathcal{A} = 1$)	83.06***	81.07***	90.56***	88.13***	78.29***	80.32***
	(9.072)	(8.900)	(8.386)	(8.347)	(7.672)	(7.486)
Observations	2,834	2,834	2,834	2,834	2,834	2,834
Plot area quintiles	No	Yes	No	Yes	No	Yes
Fixed effects	none	none	village	village	household	household

Notes: Robust-standard errors clustered on household in parentheses. Dependent variable is plot value in thousands of Rs. per acre. Constant term and area quintile dummy coefficients not reported.

The village fixed effects estimator purges locational factors correlated with both land values and active status at the village level. That the coefficient on active status does not fall (it actually rises a bit) in moving from OLS to village fixed effects indicates that these unobserved location characteristics are not a serious confound. Similarly, the finding that the coefficient on active status changes little in moving from village to household fixed effects suggests that wealth or liquidity constraints, insofar as they determine active status, are not strongly correlated with plot values. The estimates center on an average marginal value of an active plot of around 80,000 Rs/acre, representing a 25% market premium over an inactive plot.

Finally, we note a threat to the validity of our household fixed effect estimator: unobserved *plot-level* characteristics (e.g., soil quality) correlated with both land values and active status. Recall, however, that location-specific unobservables, a far more important component of residual variation across villages than across household plots *within* villages, have little impact on our regression results. This finding suggests that any bias due to unobserved plot-level characteristics is likely to be negligible.

E A Monte Carlo investigation

To evaluate the performance of our structural estimation algorithm, we conduct a Monte Carlo investigation based on 100 simulated data sets each of 2,862 plots ($\times$ 5 periods), equal to the total number of plots in our 14 map-villages. To ease computational burden, however, we only simulate an equilibrium for one of the map-villages.

Version I of the experiment takes as the true parameter values those of Table 2 in the main paper. Table E.1 reports these true parameter values (column 1) alongside, respectively, the mean, standard deviation, bias, and root mean squared error (RMSE) of the estimates across the 100 replications. Overall, our estimator performs extremely well, with all parameter estimates tightly centered around their true values.

In version II of the Monte Carlo, we lower by 10% the value of θ and increase by 10% the value of $Pr(\mathcal{D} = 1)$, with the remaining true parameter values unchanged. As reported in Table E.2, the estimator continues to perform well, showing very small bias on average along with small RMSE.

Taken together, these findings indicate that our estimation algorithm is reliable and has good small sample properties under varying conditions.

Table E.1: Monte Carlo Version I

	True	Mean	SD	Bias	RMSE
θ	16.994	16.920	0.316	-0.074	0.323
α	0.776	0.780	0.026	0.005	0.027
δ	0.722	0.712	0.033	-0.009	0.034
σ	0.648	0.638	0.020	-0.010	0.022
$Pr(\mathcal{D} = 1)$	0.670	0.671	0.029	0.001	0.029

Table E.2: Monte Carlo Version II

	True	Mean	SD	Bias	RMSE
θ	15.295	15.288	0.343	-0.006	0.341
α	0.776	0.775	0.021	-0.001	0.021
δ	0.722	0.736	0.034	0.015	0.037
σ	0.648	0.615	0.019	-0.033	0.038
$Pr(\mathcal{D} = 1)$	0.737	0.726	0.029	-0.011	0.031

F First-stage estimation details

F.1 Data

Borewell failure panel To estimate the annual probability of well failure, we use the adjacency survey to construct a 2012-16 panel of reference plot borewells that are *at risk* of failure. Wells enter the failure panel in the year after they are sunk; we drop those sunk in 2016 because they would not be at risk of failure until 2017. Since failure is an absorbing state, a well exits the panel in the year following its failure. For reasons to be discussed, when the reference plot has multiple functioning wells we only include in the panel the oldest, i.e., the first well sunk. The result is an unbalanced failure panel of 697 borewells over a maximum of five years; 320 of the 1,057 adjacencies do not contribute to the panel as they have no functioning borewells on the reference plot over the sample period. Of the 606 borewells that were functional going into 2012, about a third (195) had failed by 2016 leading to an average annual failure rate of 7.3% (see Table F.1).

Table F.1: Well Failure by Year

Year	Functional	Failed	Total
2012	559	47	606
	(92.2)	(7.8)	(100)
2013	556	41	597
	(93.1)	(6.9)	(100)
2014	527	53	580
	(90.9)	(9.1)	(100)
2015	512	33	545
	(93.9)	(6.1)	(100)
2016	489	34	523
	(93.5)	(6.5)	(100)
Total	2,643	208	2,851
	(92.7)	(7.3)	(100)

Notes: Percent of yearly total in parentheses. Sample consists of reference plot borewells subject to failure in each year.

A key issue in modelling well failure is duration dependence, as the probability of failure depends on the age of the well. If water tables were trending downward, then older and thus shallower wells would dry up first. With a non-constant hazard rate of well failure, farmers would profitably take into account not only the number of adjacent functioning borewells but also their ages, increasing the state space and thus introducing considerable complexity into the structural model. While Figure 1 suggests that water tables in our setting have been fairly stable over the last two decades, we assess the importance of duration dependence by focusing on extant borewells in 2012, when they had a median age of 12 years.

A simple test of duration dependence in well failure that avoids the intricate specification issues of duration modelling is to check whether the probability of failure between 2012-16 is related to well age in 2012, which is predetermined. The results in Table F.2 indicate significant duration dependence. The marginal effect from the column 1 estimates implies that a well that was 10 years older in 2012 has a failure rate 0.092 higher over the subsequent five years. All of this effect, however, appears to be concentrated among the 59 wells, fewer than 10% of the sample, that were more than 20 years old in 2012 (see, especially, column 3). For investment planning purposes, then, and given discounting, it is reasonable to assume that farmers view the well failure hazard as essentially constant.

Table F.2: Well Age and Subsequent Failure

	(1)	(2)	(3)
Age in 2012	0.0428***	—	—
	(0.0130)		
Age$\times\mathbb{1}_{Age\leq 10}$	—	0.00545	—
		(0.0319)	
(Age-10)$\times\mathbb{1}_{10<Age\leq 20}$	—	0.0375	—
		(0.0372)	
Age$\times\mathbb{1}_{Age\leq 20}$	—	—	0.0165
			(0.0169)
(Age-20)$\times\mathbb{1}_{20<Age}$	—	0.125***	0.132***
		(0.0456)	(0.0466)
Observations	606	606	606
log-likelihood	-375.0	-375.3	-375.5
Equal slopes test (p-value)	–	0.028	0.006

Notes: Standard errors in parentheses (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$). Dependent variable is indicator for whether well failed between 2012-16. Estimation is by ML logit. Constant term not reported. Test of equal slopes compares spline coefficients (3 in column 2 and 2 in column 3).

Borewell flow panel Data on discharge (well flow) were collected in the 2010 survey from all functioning borewells for the 2009-10 rabi season and, in the 2017 survey, for the 2016-17 rabi season. Farmers were asked to assess flow at both the beginning and end of the rabi season based on the fraction of the outlet pipe that was full when pumping water (see Giné and Jacoby 2020). The flow measure thus varies between ("minimal" coded as) 0.1 and ("full" coded as) 1.0, with one-quarter, one-half, and three-quarters flow in between. Reflecting the cyclicality of water tables during the rabi season discussed in Section 2, average flow assessments in 2016-17 (2009-10) fall from 0.84 (0.62) at the start of rabi to 0.57 (0.35) at the end. We focus here on end-of-season flows, since well interference becomes more salient as the local aquifer is drawn down.

Table F.3: End-of-Season Well Flow

	Frequency (%)	
Flow	2010	2017
0.10	32	114
	(6.2)	(22.2)
0.25	57	219
	(11.1)	(42.6)
0.50	172	143
	(33.5)	(27.8)
0.75	192	35
	(37.4)	(6.8)
1.00	61	3
	(11.9)	(0.6)
Total	514	514
	(100)	(100)
Mean	0.600	0.325
Std. dev.	0.245	0.193

To estimate flow probabilities, we construct a balanced 2-year panel of 514 functioning wells present in both the 2010 and 2017 household/plot surveys. We find that average end-of-season flow declined in the panel by almost half from 2010 to 2017 (see Table F.3), which may be attributable to differences in the respective monsoons. Rainfall in our study areas during the 2016 monsoon season (responsible for rabi 2016-2017 recharge) fell roughly 30% short of that in 2009 (see Figure A.1 for details.)

F.2 Econometric issues

Unobserved heterogeneity We allow for time invariant unobserved heterogeneity by specifying the probability of an outcome as a function of an index z_{it} for plot i in period (year) t as

$$z_{it} = \nu_i + \beta_0 + \beta_1 N_{it} + \beta_2 R_{mt} + \beta_3 a_i + \varepsilon_{it}, \tag{F.1}$$

Figure F.1: Monsoon rainfall at mandal level by year

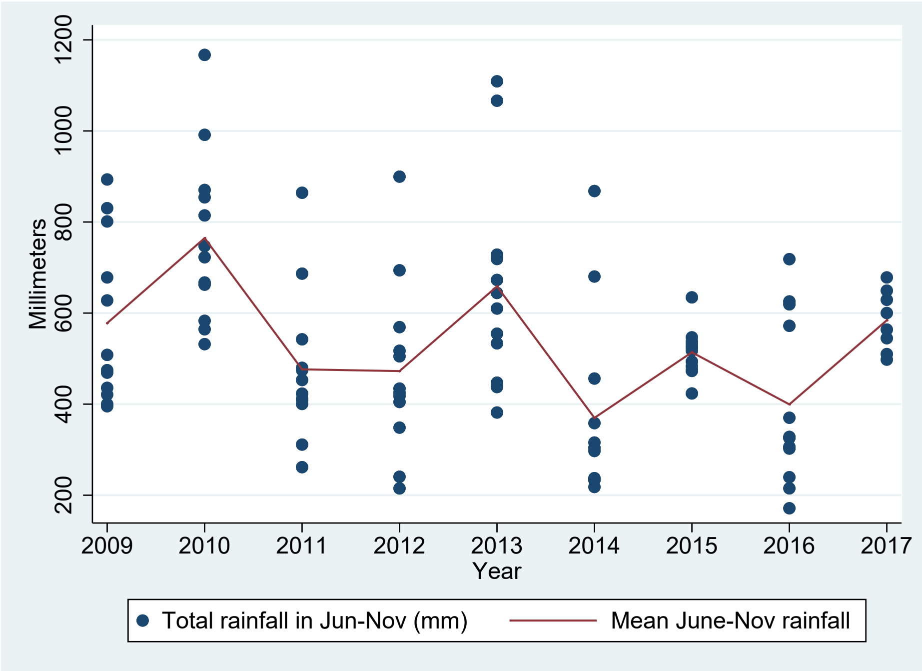

Source: For mandals in AP, the Andhra Pradesh Development Planning Society under the Planning Department of the Government of AP. For mandals in Telangana, the Telangana State Development Planning Society under the Planning Department of the Government of Telangana.

where ν_i is a random effect, N_{it} is the number of functioning wells in the adjacency at the beginning of the period, R_{mt} is a dummy variable that takes a value of one if monsoon rainfall in mandal m in year t exceeds the 2009-17 average for the 12-mandal study area as a whole, and a_i is the area of plot i. The time-varying error ε_{it} is assumed iid logistic.

A key concern is that ν_i may be correlated with (N_{it}, R_{mt}). Since nonlinear probability models do not lend themselves to fixed effects approaches (except in some special cases), we employ *correlated random effects* (CRE) under the assumption that N_{it} and R_{mt} are strictly exogenous conditional on ν_i (see, e.g., Wooldridge 2010). In particular, let

$$\nu_i \quad = \quad \gamma_1 \bar{N}_i + \gamma_2 \bar{R}_m + \mu_i, \tag{F.2}$$

where bars denote reference plot-specific means of the corresponding explanatory variable and μ_i is a continuously distributed mean zero random effect. Substituting into (F.1) yields

$$z_{it} = \mu_i + \beta_0 + \beta_1 N_{it} + \beta_2 R_{mt} + \beta_3 a_i + \gamma_1 \bar{N}_i + \gamma_2 \bar{R}_m + \varepsilon_{it}, \tag{F.3}$$

which is the index function that we use in our estimations below.

Selection bias Our estimation sample for the flow/failure model is selected on the basis of whether at least one borewell was successfully sunk on the plot prior to or (in the case of failure) during the panel period. Since the presence of a functioning borewell on a plot could be correlated with the unobserved heterogeneity driving borewell flow and/or failure on that plot, selection bias is a concern. However, insofar as ε_{it} is iid over time, and given that entry into the failure panel begins at least a "year" after the borewell is sunk (see Figure 2), this sample selection is strictly exogeneous; in other words, conditional on ν_i, selection is uncorrelated with idiosyncratic flow/failure shocks over the panel period. In these circumstances, the nonlinear CRE estimator is robust to selection bias (Wooldridge 2019).

Misclassification error It is plausible to expect adjacency survey respondents to recall the functioning status of borewells on their neighbors' plots less accurately than those on their own (reference) plots. We thus allow $\mathcal{N}_{it}$, but not n_{it}, to be subject to recall error.[1] To handle this, we assume, also quite plausibly, that the number of *existing* neighboring wells $\mathcal{N}_{it}^E$ is accurately observed. We want to estimate how the probability of a well failure at the start of rabi season of year $t+1$, which we denote by the binary indicator F_{it+1}, depends upon the true number of functioning wells in the adjacency, i.e., $Pr(F_{it+1}|\mathcal{N}_{it})$. Although the true $\mathcal{N}_{it}$ is not observed, we know that

$$Pr(F_{it+1}|\mathcal{N}_{it}^E) = \sum_{k=1}^{\mathcal{N}_{it}^E} Pr(F_{it+1}|k) Pr(k|\mathcal{N}_{it}^E), \tag{F.4}$$

where $Pr(k|\mathcal{N}_{it}^E)$ is the discrete probability density of the true number of functioning wells outside of the reference plot. This density is of the binomial form

$$Pr(k|\mathcal{N}_{it}^E) = \frac{\mathcal{N}_{it}^E!}{k!(\mathcal{N}_{it}^E - k)!} p^{\mathcal{N}_{it}^E - k}(1-p)^k, \tag{F.5}$$

where p is the underlying annual probability of well failure.[2] The misclassification error model (MEM) estimator then assumes that the likelihood contribution conditional on

[1] While there has been recent progress in the econometrics literature on models of misclassification (e.g., Mahajan 2006; Hu 2008), no tractable general approaches exist applicable to our specific situation.

[2] More precisely, p should be thought of the average failure rate of borewells in the adjacency of reference plot i, excluding those on the reference plot itself. We show empirically below that p is a function of the number of functioning borewells in the relevant neighborhood. In the context of equation (F.4), the relevant neighborhood is that around each of the plots *adjacent* to the reference plot, about which we have no adjacency-level data. Given this, we approximate p using the district average of the actual failure rate, yielding $\hat{p} = 0.104$ for Anantapur and $\hat{p} = 0.052$ for Mahabubnagar.

unobservable μ is

$$\ell_i^y(\mu) = \prod_{t=1}^{T} \sum_{k=1}^{\mathcal{N}_{it}^E} Pr(F_{it+1}|z_{it}(k,\mu))Pr(k|\mathcal{N}_{it}^E, \hat{p}). \tag{F.6}$$

F.3 A joint model of well flow and failure

The well flow and failure panels cover 382 and 697 adjacencies, respectively, of which 360 overlap, i.e., include wells with both flow and failure observations.[3] This overlap allows identification of the correlation in reference plot-level unobserved heterogeneity between the well flow and failure processes. Such correlation is plausible if well failure is seen as a state of zero flow forever.

We discuss the likelihood contribution of each process in turn and then derive the joint flow-failure likelihood.

Flow To estimate the probabilities for the five well-flow states ($q = 0.1, 0.25, 0.5, 0.75, 1.0$), we use a CRE ordered logit for the two-year panel. The conditional likelihood contribution of reference plot i is

$$\ell_i^f(\mu) = \prod_t \prod_{m=1}^{5} \left(\frac{1}{1 + e^{c_{m+1} + z_{it}^f(\mu)}} - \frac{1}{1 + e^{c_m + z_{it}^f(\mu)}} \right)^{\mathbb{1}_{Q_{it}=m}} \tag{F.7}$$

where $z_{it}^f(\mu)$ is a linear index for flow as in equation (F.3), Q_{it} is a 5-valued flow-state indicator and the c_m are cutoff parameters with $c_1 = -\infty$ and $c_6 = \infty$.[4]

Failure For reasons already noted, we adopt a constant failure hazard specification, using the sequential logit as in Cameron and Heckman (1998), among others. The condi-

[3] Non-overlap occurs because flow data were collected on all borewells owned by the household, irrespective of their inclusion in the adjacency survey, and because there are adjacencies that did not have functioning wells on the reference plot in 2010 and 2017, when flow data were collected.

[4] Our measure of the number of neighboring borewells differs between flow and failure estimation datasets. In the former case, owners of functioning wells were asked about the number of functioning borewells within a 100 meters radius of the reference plot, which is not precisely the same as the number in the adjacency. In practice, however, the total number of borewells within 100 meters averages 2.40 as compared to 2.36 in the average agency (for the 588 reference plots with both measures available). Since the $\mathcal{N}_{it}$ used in the flow estimation is the contemporaneous (rather than retrospective) report of the respondent, we assume no misclassification error in (F.7).

tional likelihood contribution with $\mathcal{N}_{it}$ subject to misclassification error is

$$\ell_i^F(\xi) = \prod_{t=\tau_i-1}^{T_i-1} \sum_{k=1}^{\mathcal{N}_{it}^E} \frac{e^{z_{it}^F(k,\xi)\cdot F_{it+1}}}{1+e^{z_{it}^F(k,\xi)}} \cdot Pr(k|\mathcal{N}_{it}^E, \hat{p}), \tag{F.8}$$

where $z_{it}^F(\xi)$ is a linear index for failure, F_{it+1} is the binary failure indicator defined above, τ_i is the year that the borewell first enters the panel (or 2012, whichever comes last), T_i is the last year the borewell exists in the panel (or 2016, whichever comes first), and ξ is the unobserved heterogeneity in well failure. For reference plots with multiple wells, only the first one sunk is included in the failure panel. Allowing multiple borewells on a plot would lead to a violation of strict exogeneity due to correlation between N_{it} and the failure shock.

The joint model For the joint flow/failure estimation, we follow, e.g., Eckstein and Wolpin (1999) in assuming that the reference plot level random effects, μ and ξ, are linearly related, i.e., $\xi = \kappa\mu$, where κ is a covariance parameter. Defining three indicator variables, D_i^1, D_i^2, D_i^3 for whether reference plot i contributes, respectively, only flow data, only failure data, or both flow and failure data, and assuming that μ has a discrete distribution with J points of support $(\mu_1, ..., \mu_J)$ and associated probabilities $(\rho_{\mu 1}, ..., \rho_{\mu J})$, the full log-likelihood is

$$\mathcal{L} = \sum_i \log \left\{ \sum_j \rho_{\mu j} \ell_i(\mu_j, \kappa\mu_j) \right\}, \tag{F.9}$$

where

$$\ell_i(\mu_j, \xi_j) = \left[\ell_i^f(\mu_j)\right]^{D_i^1} \left[\ell_i^F(\xi_j)\right]^{D_i^2} \left[\ell_i^f(\mu_j)\ell_i^F(\xi_j)\right]^{D_i^3}. \tag{F.10}$$

To estimate the probabilities of the five well flow states, $\pi_k(N, R, \nu^f)$, and the failure probability, $\pi_F(N, R, \nu^F)$, where ν^f and ν^F are, respectively, the flow and failure unobserved heterogeneity *unconditional* on the CRE covariates $(\bar{N}_i, \bar{R}_m)$, as defined in equation (F.2), we proceed in two steps:

Step 1: Maximize the CRE likelihood given by equation (F.9), including stepwise w.r.t. the number of points of support J, and obtain estimates of the linear index coefficients $\hat{\beta}^f, \hat{\gamma}^f, \hat{\beta}^F$, and $\hat{\gamma}^F$ (see equation F.3).

Step 2: Set $\beta^f = \hat{\beta}^f$, $\beta^F = \hat{\beta}^F$, $\gamma^f = \gamma^F = 0$, and re-maximize the likelihood with respect to the unconditional heterogeneity distribution parameters $(\nu_1, ..., \nu_J)$ and

$(\rho_{\nu 1}, ..., \rho_{\nu J})$, adding, as in Step 1, points of support until the likelihood fails to improve. Based on these estimates, compute $\pi_k(N, R, \nu_j^f)$ and $\pi_F(N, R, \nu_j^F)$.

The top panel of Table F.4 reports the coefficient estimates from Step 1. Column 1's specification ignores misclassification error, whereas column 2's uses the MEM approach to correct for it. Corroborating the well interference externality, we find that having more borewells in an adjacency depresses flow and makes failure of the reference well more likely, but only when misclassification error is taken into account. Also, having had a "good" previous monsoon (above mean rainfall) improves well flow but does not have a significant effect on failure, consistent with our assumption that well failure is an absorbing state, independent of the vagaries of the monsoon. Borewells located in larger plots have both significantly better flow, conditional on the number of functioning wells in the neighborhood, and a lower probability of failure. Moving to Step 2, we find that two discrete types, with associated probabilites reported in the second panel of Table F.4, fit the data adequately inasmuch as adding a third type leads to a virtually identical Step 2 likelihood value. The negative cross-equation error correlation indicates that a borewell with a good flow, ceteris paribus, is also less likely to fail.

Robustness Instead of a random effects ordered logit for the five flow states, we now run a linear regression with reference plot fixed effects, where the dependent variable takes on values from 1 to 5 (col. 1 of Table F.5). Since $\log(N)$ in this case is not subject to recall error, we do not instrument for it. In the case of failure, we estimate a linear probability model with reference plot fixed effects (col. 2) and by FE-IV (col. 3) using $\log(\mathcal{N}^E + n)$ as an instrument for $\log(\mathcal{N} + n)$. Estimation samples for the separate flow and failure models are identical to those used in the joint nonlinear estimation as reported in Table F.4. Results for the two sets of procedures are qualitatively similar.

Table F.4: Joint Flow-Failure model

Equation/covariate	(1)	(2)
Flow:		
log(N)	-0.899	-0.899
	(0.172)	(0.172)
Good monsoon	1.908	1.908
	(0.803)	(0.803)
Log plot area	0.229	0.229
	(0.0716)	(0.0716)
Failure:		
log(N)	0.0632	1.140
	(0.599)	(0.472)
Good monsoon	-0.255	-0.259
	(0.259)	(0.259)
Log plot area	-0.324	-0.327
	(0.100)	(0.100)
Estimation method	CRE	CRE-MEM
Log-likelihood	-2168	-2165
Heterogeneity (ν)		
Low type probability	0.0191	0.346
	(0.00963)	(0.0964)
High type probability	0.981	0.654
	(0.00963)	(0.0964)
Standard deviation	2.233	0.229
	(73.73)	(0.149)
Cross-equation correlation	-0.0950	-0.0699
	(1.249)	(0.0449)

Notes: Standard errors in parentheses. Maximum likelihood estimates with reference plot-level correlated random effects (CRE). Ordered logit cutoffs for flow, constant term for failure, mandal dummy and CRE covariate coefficients for both equations not reported. Sample size = 3,401. Standard deviation refers to that of ν^f, the unobserved heterogeneity unconditional on the CRE covariates; cross-equation correlation is $corr(\nu^f + \varepsilon^f, \nu^F + \varepsilon^F)$.

Table F.5: Determinants of well flow and failure–linear models

	flow (1-5)	failure (0/1)	
	(1)	(2)	(3)
log(N)	-0.513***	0.00368	0.283***
	(0.106)	(0.0523)	(0.0509)
Good monsoon	1.110***	0.00881	0.00749
	(0.347)	(0.0118)	(0.0129)
Reference plot FE	YES	YES	YES
Observations	1,028	2,851	2,851
Number of ref. plots	514	697	697

Notes: Standard errors in parentheses clustered by reference plot (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$). Columns 1 and 2 are by ordinary least-squares; columns 3 is by two stage least squares using the log number of existing wells in adjacency as an instrument.

G Heterogeneity in land productivity

In this section, we evaluate whether our model adequately accounts for unobserved heterogeneity in land productivity. As discussed in Appendix Section D, we have data on self-assessed plot values, which presumably reflect their respective owners' expectations over future agricultural productivity. A more productive plot, in the eyes of its owner, will attract more drilling and thus tend to have more functioning borewells. This observation suggests using the correlation between plot value and the number of functioning borewells on the plot as an external validation target; specifically, we consider the slope coefficient from the regression of value per acre V_i on the average number of functioning wells $\bar{n}_i$ in plot i over the previous five years:

$$V_i = a + b\bar{n}_i + e_i, \tag{G.1}$$

using only active plots (i.e., $\mathcal{A}_i = 1$).

Our statistical test depends on $\hat{b}$ increasing in the variance of the land productivity parameter θ. However, since $\hat{b} = cov(V_i, \bar{n}_i)/var(\bar{n}_i)$, $\hat{b}$ is not unambiguously increasing in $var(\theta)$; both numerator and denominator can vary. To check how $\hat{b}$ moves in practice, we therefore run the following exercise: First, using simulated data from our baseline model, we obtain $\hat{b} = 194$ ('000 Rs/acre; see Table G.1, row 1). Next, we introduce variance in the i.i.d. (across plots) productivity term θ using simple two-point distributions with successively larger spreads, re-estimating the other model parameters accordingly. Simulating data from these new models, we find that the resulting values of $\hat{b}$ all exceed 194 (Table G.1, rows 2 and 3). In sum, additional $\theta-$heterogeneity would only increase $\hat{b}$.

Turning to the actual data, we find that $\hat{b} = 95$. A formal one-sided test cannot reject the null hypothesis that the model-generated $\hat{b}$ (i.e., 194) is equal to this empirical coefficient against the alternative that the model under-shoots $\hat{b}$ because it does not fully account for unobserved heterogeneity (in fact, our model over-shoots). The p-value for this bootstrap test (Hansen 2022) with 10,000 replications is 1.000.[5]

In sum, the sources of heterogeneity built into the baseline model are sufficient, indeed more than adequate, to rationalize the observed link between land value and borewell numbers. Adding $\theta-$heterogeneity would only worsen model fit to this moment.[6]

[5]This test is robust to classical measurement error in plot values because measurement error in a dependent variable does not bias estimated regression coefficients. By contrast, an alternative validation strategy that tries to match the standard deviation of plot values would conflate measurement error with true heterogeneity in land productivity.

[6]Strictly speaking, in the actual data, land values reflect the present discounted stream of profits from both irrigated and rain-fed cultivation, whereas land values generated from the model reflect only the incremental profit from irrigated agriculture. Moreover, the non-irrigated component of actual land

Table G.1: Evaluating $\theta-$heterogeneity

	$\hat{b}$
Model:	
$CV(\theta) = 0.0$ (baseline)	194
$CV(\theta) = 0.1$	200
$CV(\theta) = 0.2$	223
Data:	95
	[64,125]

Notes: Percentile-t confidence interval, with clustering at household level, in square brackets (sample size = 1,098 plots owned by 907 households). $CV(\theta)$ is the coefficient of variation of TFP parameter θ. Models with $CV(\theta) > 0$ assume a discrete two-point, symmetric, distribution around the baseline model estimate of θ.

values may also be (positively) correlated with θ. In this case, however, the $\hat{b}$ from the actual data would *exceed* the one that should be compared to our model-generated $\hat{b}$, thus only strengthening our conclusion that there is enough heterogeneity in the model already to rationalize the data.

H Drilling: Strategic substitutability

Using a five-year panel (2012-16) on 1,057 reference plots covered by the adjacency survey, we estimate a linear probability model for drilling of the form

$$d_{it} = \alpha_i + \beta_1 \mathcal{N}_{it} + \beta_2 R_{mt-1} + \varepsilon_{it}, \tag{H.1}$$

where α_i is a reference plot fixed effect. We assume classical measurement error in $\mathcal{N}_{it}$ and use the number of existing wells in the adjacency (outside the reference plot), $\mathcal{N}_{it}^E$, as an instrument.

Identification While the strategic substitutability parameter β_1 resembles a peer effect, causal identification is not as challenging as implied by Manski (1993). In particular, since $\mathcal{N}_{it}$ reflects past (as opposed to current year) drilling by neighbors, and we are taking out plot fixed effects, β_1 is identified even if contemporaneous plot-specific drilling shocks ε_{it} are (spatially) correlated between plots in the same adjacency.[7] While identification does break down if ε_{it} is both spatially *and* serially correlated, a spurious finding of strategic substitutability could only be explained by either negative spatial or negative serial correlation in drilling shocks, either of which is implausible. In the likelier scenario of positive spatial and serial correlation of drilling shocks, our estimate of β_1 would be biased upward, i.e., toward zero, and thus *away* from strategic substitutability.[8] Indeed, we test for one such source of bias below by conditioning on the number of own borewells on the reference plot.

Results Column (1) of Table H.1 reports fixed effects least-squares estimates showing zero impact of neighboring wells on drilling. Column (2) displays the first stage regression of $\mathcal{N}_{it}$ on the instrument $\mathcal{N}_{it}^E$ and column (3) the resulting FE-IV estimate. We find a significantly negative effect of neighboring wells once we instrument for measurement error. One concern, noted above, is that, if there is spatial correlation in the unobservables, then $\mathcal{N}_{it}^E$ may be correlated with the residuals, which contain the effect of *own* borewells on drilling. To assess this, in column (4) we add dummies for the number of borewells

[7] Pfeiffer and Lin (2012) estimates the effect of a neighbors' groundwater pumping on simultaneous own pumping behavior using a cross-sectional instrumental variables strategy.

[8] A referee points out that this argument hinges on the nature of the common property externality. If groundwater extraction involves an important stock externality (see fn. 2 in the main text), then there could be strategic complementarity in drilling decisions, wherein one farmer's attempt to capture the resource encourages other potential users to drill. In this case, positive spatial and serial correlation of drilling shocks would bias us *toward* strategic complementarity. As noted, however, the stock externality is not operative in our setting.

on the reference plot to remove the effect of own borewells from the residuals.[9] That there is no appreciable difference between the estimates of β_1 across columns (3) and (4) gives us further confidence that negative effect of neighboring wells on reference plot drilling is indeed causal. Finally, in columns (5) and (6), for the purposes of validating the structural model in Section 4.5, we replicate the column (3) and (4) specifications, respectively, dropping observations with more than one functioning well on the reference plot.

Table H.1: Determinants of drilling 2012-16–Linear probability models

	(1) drill	(2) $\mathcal{N}$	(3) drill	(4) drill	(5) drill	(6) drill
No. func wells exc ref plot ($\mathcal{N}$)	-0.0154		-0.0482**	-0.0485**	-0.0442**	-0.0381**
	(0.0108)		(0.0213)	(0.0199)	(0.0198)	(0.0192)
Good Monsoon (R)	-0.00340	0.000760	-0.00248	-0.00257	0.00124	0.000391
	(0.00881)	(0.00775)	(0.00884)	(0.00866)	(0.00888)	(0.00877)
No. exist wells exc ref plot		0.900***				
		(0.0285)				
1 func well on ref plot				-0.243***		-0.256***
				(0.0241)		(0.0241)
2 func wells on ref plot				-0.447***		
				(0.0538)		
Reference plot FE	YES	YES	YES	YES	YES	YES
Observations	5,285	5,285	5,285	5,285	4,837	4,837
Number of ref. plots	1,057	1,057	1,057	1,057	988	988

Notes: Standard errors in parentheses clustered by reference plot (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$). Columns 1 and 2 are by ordinary least-squares; columns 3-6 are by two stage least squares using the number of existing wells in adjacency (outside of reference plot) as instrument. Columns 5 and 6 drop observations with more than one functioning well on the reference plot.

[9] Insofar as past drilling successes lead to more borewells on the reference plot, the fixed effects estimator of the own borewell coefficients in this short panel are biased (as per Nickell 1981). Thus, we treat the column (4) results as a specification test.

I Transitional dynamics

I.1 Equilibrium

We now describe the Adjacency Equilibrium over the transition path of the benchmark village economy to a new steady state following the introduction of a tax τ on borewells at date $t = 1$. The village map transits to a new (steady-state) Adjacency Equilibrium AE_τ. We assume that i) CCPs along the transition depend only on the state of the adjacency and date t, and ii) that the plot owner has beliefs about the evolution of $X_{(i1)t}$ along an "average" transition. Assumption ii) requires that equilibrium beliefs about the state of the adjacency at date t be correct when averaged over the map's stochastic transition paths. Thus, we have

Definition: Let $F_0^\infty(X)$ (or F_0^∞ in short) be the stationary distribution over the state of the map at the initial Adjacency Equilibrium AE_0. An Adjacency Equilibrium over the transition path is a vector of choice probability functions $\{CCP_{it}(X_{(i1t)}, n_{it})\}_{t=1}^\infty$ and of beliefs $\{\tilde{F}_{it}^*(X_{(i1)t+1}|X_{(i1t)}, n_{it})\}_{t=1}^\infty$ such that: a) CCPs and beliefs converge to the CCPs and beliefs of AE_τ; b) given beliefs $\tilde{F}_{it}^*$, the decision rule CCP_{it}^* is the solution of plot owner i's dynamic game "against nature" at every t; and c) beliefs at each t are correct on "average". That is, let $F_t(X_t; \{CCP_s\}_{s=1}^t, F_0^\infty)$ be the joint distribution over the state induced by the primitives, the vector of CCPs from date $s = 1$ to t and the initial steady state distribution of the map F_0^∞. Further, let $F_t(X_{(i2)t}|X_{(i1)t}, n_{it}\{CCP_s\}_{s=1}^t, F_0^\infty)$ be the conditional distribution implied by $F_t(X_t; \{CCP_s\}_{s=1}^t, F_0^\infty)$. Then,

$$\tilde{F}_{it}^*(X_{(i1)t+1} = x_{(i1)t+1}|X_{(i1t)} = x_{(i1)t}, n_{it}) = \sum_{x_{(i2)t}} F_t(x_{(i2)t}|x_{(i1)t}, n_{it}; \{CCP_s\}_{s=1}^t, F_0^\infty)$$
$$F_t(x_{(i1)_{t+1}}|x_{(i1)t}, x_{(i2)t}, n_{it}; CCP_t). \quad \text{(I.1)}$$

I.2 Solution algorithm

Recall that in our empirical structural model we reduce the dimensionality of the AE by partitioning the set of adjacencies into types, such that all adjacencies of the same type share beliefs and CCPs, and we limit the state of the adjacency to $(\mathcal{N}, n)$. We compute an AE along the transition path as follows:

Step 0 Solve for the steady state in the benchmark no-tax economy ($\tau = 0$) using the algorithm in the main text and recover the steady state distribution F_0^∞.

Step 1 Solve for the steady state in the counterfactual economy ($\tau > 0$) using the algorithm in the main text and recover the value function for each plot type (V_T).

Step 2 Assume that the village converges to this counterfactual steady state and that it is in this steady state in period T.

Step 3 Guess a sequence of beliefs $\{\tilde{F}_t\}_{t=1}^T$ (as an initial guess, linearly interpolate beliefs from the benchmark to the counterfactual steady state).

Step 4 Solve for plot owner's decision as follows:

Step 4.1 Start in period T-1.

Step 4.2 Given value function V_T and beliefs $\tilde{F}_{T-1}$, solve for the CCP_{T-1} and recover V_{T-1}.

Step 4.2 Iterate until $t = 1$ and recover $\{CCP_t\}_{t=1}^T$.

Step 5 Given $\{CCP_t\}_{t=1}^T$, sample an initial state of the map from the benchmark economy in the steady state F_0^∞ and simulate a transition of well drilling decisions, successes, and failures in every plot on the map from $t = 1$ to T. Replicate this simulation N_S times, e.g. $N_S = 250$.

Step 6 From the transition simulations, construct estimates of the one-period ahead state transition matrices $F_t(\mathcal{N}'|\mathcal{N}, n)$ for each plot type (i.e., averaging across plots on the map of the same type at the same date). Update beliefs and go back to Step 4 and continue iterating through step 6.

Step 7 Stop once the incremental change in $\{CCP_t\}_{t=1}^T$ is sufficiently small.

J Spatial distribution of drilling

We simulate steady states in the network and island economies for one of our map-villages, Ayyavaripalli, with the θ parameter calibrated in the island economy so that the resulting borewell density matches that of the network economy in this village. We then simulate a long history of drilling decisions by farmers, along with well failures, and compute average drilling rates at the plot level.

Figure J.1 shows the spatial distribution of average drilling rates under both the network and island economies, along with the corresponding histograms. In both maps, colors represent deciles of drilling probability, with yellow indicating plots with the highest drilling rates and darker shades indicating lower rates. In the network economy (left), farmers must drill more frequently to sustain well density due to local externalities. In the island economy (right), fewer drilling attempts are needed, as wells are less likely to fail given the assumed absence of externalities.

Figure J.1: Distribution of drilling rates in Ayyavaripalli

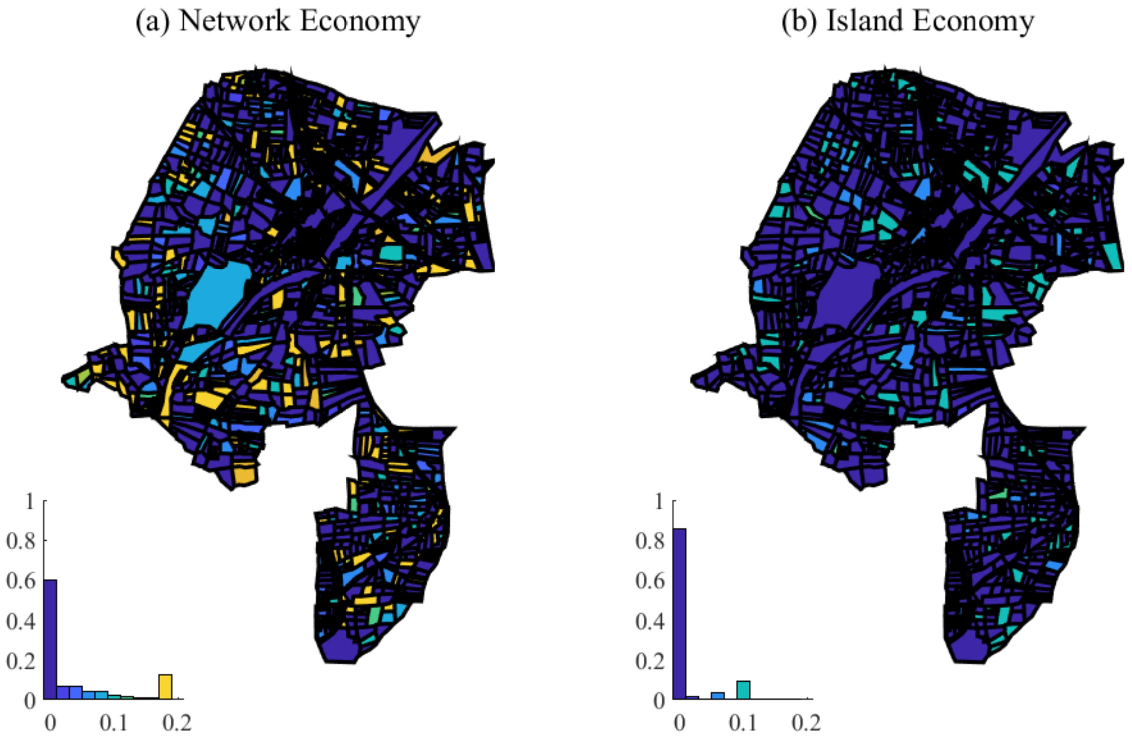

References

Cameron, Stephen V and James J Heckman, "Life cycle schooling and dynamic selection bias: Models and evidence for five cohorts of American males," *Journal of Political economy*, 1998, *106* (2), 262–333.

Eckstein, Zvi and Kenneth I Wolpin, "Why youths drop out of high school: The impact of preferences, opportunities, and abilities," *Econometrica*, 1999, *67* (6), 1295–1339.

Giné, Xavier and Hanan G Jacoby, "Contracting under uncertainty: Groundwater in South India," *Quantitative Economics*, 2020, *11* (1), 399–435.

Gulati, Mohinder and Sanjay Pahuja, *Direct delivery of power subsidy to agriculture in India*, Sustainable Energy for All, 2012.

Hansen, Bruce, *Econometrics*, Princeton University Press, 2022.

Manski, Charles F, "Identification of endogenous social effects: The reflection problem," *The review of economic studies*, 1993, *60* (3), 531–542.

Nickell, Stephen, "Biases in dynamic models with fixed effects," *Econometrica: Journal of the econometric society*, 1981, pp. 1417–1426.

Pfeiffer, L. and C. Y. C. Lin, "Groundwater pumping and spatial externalities in agriculture.," *Journal of Environmental Economics and Management*, 2012, *64* (1), 16–30.

Wooldridge, Jeffrey M, *Econometric analysis of cross section and panel data* 2010.

_ , "Correlated random effects models with unbalanced panels," *Journal of Econometrics*, 2019, *211* (1), 137–150.