Mathematical Modeling and Algorithm Application
ISSN: 3006-0842 | Vol. 10, No. 1, 2026

A Compartmental Differential-Equation Model of the Imperial Examination and Social-Class Mobility in Ming China

Zihan Wu*

The Experimental High School Attached to Beijing Normal University, Beijing 100032, China
* **Corresponding author**: Zihan Wu (Email: wzh_0156@163.com)

Abstract: The imperial examination (keju) is conventionally credited with keeping the late-imperial Chinese elite open to talent, yet that openness is usually measured as a rate for a period or a cohort rather than modelled as a process. This paper builds a compartmental ordinary differential equation (ODE) model in which two household stocks — commoner households C and elite households E, the latter defined as households with a degree or an office-holding within three generations — interact through an examination gate whose selection function $\pi(t)=C/(C+\rho E)$ is exactly the quantity historians report as the share of newly enrolled jinshi coming from strictly commoner families. The model is calibrated on a dataset assembled for this study from published scholarship and public databases: the per-examination jinshi intake for 88 examinations (24,588 graduates) and the per-examination origin composition of 57 examinations (15,528 usable observations). Two parameters are pinned to documented historical anchors (elite decay $\delta=0.034$ yr^{-1} from the finding that close to 87% of jinshi lineages did not survive three generations; the initial commoner stock from the 1393 census of 10,652,870 households); the remainder are estimated by weighted least squares, with class-transfer terms that are conservative by construction. The fitted model reproduces the observed trajectory with a weighted R^2 of 0.746 and a mean absolute percentage error of 7.6%. A specification in which the elite stock grows without limit fits markedly worse (R^2=0.609), and the two specifications are separated on fit quality; the implied elite stock of the retained specification grows 14- to 20-fold over the dynasty (14.5-fold at the fitted point; the level of the stock and the elite advantage ρ are identified only jointly), against an independently documented 16.7-fold rise in the licentiate population that played no part in the fitting. Sensitivity analysis, computed as partial rank correlation coefficients, ranks elite decay first (PRCC +0.97), the degree quota second (−0.95) and the entry barrier facing commoners a distant third (−0.09, marginal at p=0.049). The jinshi quota is almost inert: doubling it moves the endpoint commoner share by about 0.2 percentage points, because the metropolitan channel is far narrower than the elite stock it feeds.

Keywords: Imperial Examination; Social Mobility; Compartmental Model; Elite Reproduction; Parameter Estimation; Ming China

1. Introduction

The civil examination system shaped the Chinese imperial state for more than a millennium, and its social consequences have been debated for nearly as long. The classical quantitative statement of the problem is Ho Ping-ti's The Ladder of Success in Imperial China [1], which coded the three-generation pedigree recorded in the dengke lu (examination registers) for thousands of jinshi and concluded that roughly half of Ming graduates came from families with no office-holding in the preceding three generations. That finding has been re-examined repeatedly. Xu Hong re-coded a substantially larger sample and confirmed the direction of Ho's result while reporting a lower commoner share for the late Ming [2]; Guo Peigui has produced the most extensive counts of degree-holders and of commoner-origin proportions [3], and has documented the size of the provincial juren population [4]; and a parallel econometric literature has exploited the examination system's geography for identification, using prefecture-level jinshi density to study persistence [5] and the abolition of the examination to study political stability [6]. The measurement of mobility as such has its own methodological tradition, running from Prais's formulation of mobility indices [7] to the Markov-chain models of social processes [8], and the debate over how open the Ming examination really was remains live [9].

These literatures share a common feature: they treat mobility as a statistical quantity to be measured rather than as a process to be modelled. Rates are computed for periods or cohorts, compared across sources, and interpreted; but the feedback structure that generates them — the accumulation of elite households, their subsequent decay, and the gate that regulates entry — is left implicit. The result is that the degree of openness can be stated, but not what determined it, nor which institutional lever would have changed it. Where dynamical models have been brought to bear on pre-modern stratification, they have addressed collapse and instability rather than mobility: the elite–commoner system of the HANDY model [10], the demographic–fiscal models of secular cycles [11,12] and the elite-overproduction mechanism in structural-demographic theory [13] all describe what happens when elites accumulate, but they take the recruitment channel as given.

This paper addresses that gap with a deliberately small dynamical model. The examination system is treated as a gate between two household stocks, commoner households $C(t)$ and elite households $E(t)$, and the reported share of new jinshi coming from strictly commoner families, $\pi(t)$, is the observable generated by that gate. The model is intentionally minimal: two state equations, an exogenous examination quota taken from the historical record, and one selection function with a single behavioural parameter. Its purpose is

not to forecast but to expose the mechanism and, in particular, to test whether the observed trajectory requires anything beyond elite accumulation and decay. The compartmental form follows the standard susceptible–infectious construction of Kermack and McKendrick [14], with households rather than individuals as the unit. Two features of the data shape the analysis. The observable must match the definition of E exactly, which is why the strictly commoner share, rather than the broader commoner-origin share, is used and why Section 6.4 reports both conventions; and π identifies only the trajectory of the elite stock, not its level, which is why an independent licentiate series that never enters the fitting is brought in as an external check in Section 6.2.

2. Symbols and Notation

Table 1 fixes the notation. It also states the source of each quantity, distinguishing those pinned to documentary evidence from those estimated in Section 6.

Table 1 Symbols, units and sources.

Symbol	Meaning	Unit	Value / source
t	Calendar year (AD)	yr	1371–1622
$C(t)$	Commoner household stock (no degree and no office in three generations)	households	C_0 = 10,652,870 (1393 census [15,16])
$E(t)$	Elite household stock (a degree or an office within three generations)	households	latent
$m(t)$	Annual jinshi intake (exogenous)	persons yr−1	per-examination counts ÷ 3 [17]
$\pi(t)$	Strictly commoner share of new jinshi	—	observable [2]
g_C	Intrinsic growth rate of commoner households	yr−1	estimated, bounded above at 0.10 (Section 6.5)
K_C	Carrying capacity of commoner households	households	estimated
g_E	Gross expansion rate of the elite stock	yr−1	estimated
K_E	Carrying capacity of the elite stock	households	estimated
ρ	Relative examination advantage of elite households	—	estimated jointly with the level of E (Section 6.1)
δ	Annual exit rate of elite households	yr−1	0.034 (pinned, [3])
λ	Elite households created per commoner-origin jinshi	households per person	1 (assumed)
q	Degree-quota multiplier (sensitivity)	—	1.0 baseline
q_m	Jinshi-quota multiplier (sensitivity)	—	1.0 baseline
$R(t)$	Imperial-clan population (five-class extension)	persons	g_R = 0.0351 yr−1 (fitted to register counts, [15])

3. Model Assumptions

(1) Wars, epidemics and natural disasters are treated as absent from the mobility process over the period considered; their short-run effects are not modelled.

(2) The unit of state is the household, not the individual. A household does not switch class repeatedly within one generation, and a successful candidate induces at most one upward transition for his household ($\lambda = 1$).

(3) Natural growth of both household stocks is logistic and layer-homogeneous, with capacities K_C and K_E constant over the period.

(4) Examination quotas are exogenous. The state of the class structure does not feed back into the number of degrees awarded, which were fixed by statute and by provincial quota [18].

(5) The triennial examination is approximated by a continuous flow $m(t)$, taken as a step function equal to each examination's intake divided by three.

(6) Regional heterogeneity is excluded. The northern and southern registers (nanbei bang) and the province-specific provincial quotas are not modelled and are noted as a boundary of the model in Section 6.5.

4. Method

4.1. Three-class base model

Fig. 1a shows the structure. Commoner households C become elite households by producing a successful jinshi; elite households revert to commoner status when their degree capital is exhausted. The equations are

$$\dot{C} = g_C\, C\left(1 - \frac{C}{K_C}\right) - \lambda\, m(t)\, \pi(t) + \delta\, E \qquad (1)$$

$$\dot{E} = g_E\, E\left(1 - \frac{E}{K_E}\right) - \delta\, E + \lambda\, m(t)\, \pi(t) \qquad (2)$$

$$\pi(t) = \frac{C(t)}{C(t) + \rho\, E(t)} \qquad (3)$$

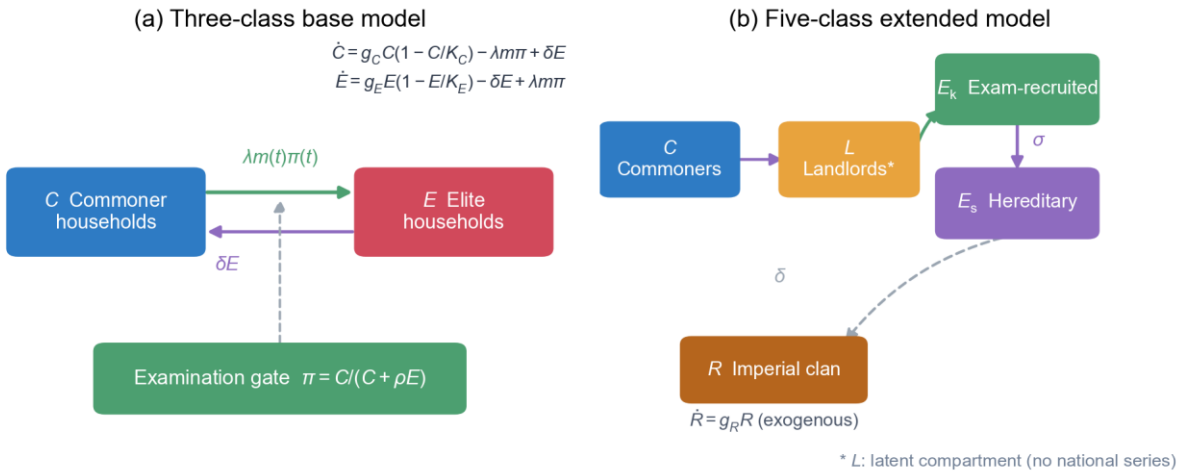


Fig. 1 Model structure. (a) Three-class base model: commoner households C and elite households E coupled through the examination gate $\pi = C/(C + \rho E)$; $\lambda m\pi$ is the upward flow, δE the downward flow. (b) Five-class extension: a latent landlord compartment L mediates the transition from commoners to examination-recruited officials E_{k}, hereditary officials E_{s} are fed by privilege at rate σ, and the imperial clan R grows exogenously.

The two transition terms, $\lambda m\pi$ and δE, appear with opposite signs in (1) and (2) — a household that leaves one stock is added to the other — so the class-transfer flow is conservative by construction: with $N = C + E$, $\dot{N} = g_C C(1 - C/K_C) + g_E E(1 - E/K_E)$ contains no transfer term. The gate equation deserves comment, because it is where the model's interpretative content lies. If households of the two classes produced successful candidates in proportion to their numbers alone, the commoner share of new jinshi would be $C/(C + E)$, and since C exceeds E by two orders of magnitude the share would be pinned near unity. The parameter ρ is therefore the model's central behavioural quantity: it is the per-household relative advantage of an elite household in the examination, capturing endowments that the historiography has long emphasised — tutors, libraries, examination craft transmitted across generations, and the leisure to study. At the fitted point, $\rho \approx 85$, elite households are a numerically negligible minority that nonetheless takes a large share of the degrees; because π depends on E only through ρE, the value of ρ and the level of E are identified only jointly (Section 6.1). The structure is exactly the one that Ho and Xu measured, written as a selection function.

Two parameters are pinned rather than estimated. The elite exit rate δ follows from the finding that close to 87% of jinshi lineages failed to reproduce their status within three generations, roughly 60 years, and that 96% did not survive beyond the third [3], giving $\delta = \ln(1/0.13)/60 \approx 0.034$ yr−1; this is the rate at which a household leaves the three-generation pedigree window. The initial commoner stock is the 1393 household census, $C_0 = 10{,}652{,}870$ [15,16]. The remaining quantities — g_C, K_C, g_E, K_E, ρ and the initial elite stock — are estimated from the observed π series by weighted least squares, with weights proportional to the square root of each examination's usable sample size; g_C is bounded above at 0.10 yr−1 for the reason given in Section 6.5.

Nondimensionalisation and identifiability. Measuring C in units of C_0 and absorbing ρ into the elite variable, $e = \rho E/C_0$, leaves the system in terms of c, e, g_C, $\kappa_C = K_C/C_0$, $\kappa_E = \rho K_E/C_0$ (the capacity of the scaled elite variable e), g_E, ρ and e_0. Because π depends on E only through the product ρE, the level of the elite stock is not identified by π alone; only its relative trajectory is. The level must be fixed by an external series; this is taken up in Section 6.2.

4.2. Five-class extension

The base model collapses the elite into one stock. The five-class specification distinguishes commoners, landlords, examination-recruited officials, hereditary officials and the imperial clan; Fig. 1b shows the structure. The three additions differ sharply in what the data can support.

The imperial clan is the one component with its own series and calibrated rate. The clan population entered in the imperial genealogy registers rose from 8,203 in 1529 to 19,611 in 1553 and 28,840 in 1565, and a memorial of 1595 put the total, including women and members not yet enrolled, at about 157,000 [15]. A log-linear fit over the three register counts, which share one definition, gives $g_R = 0.0351$ yr−1 (doubling time 19.7 years) with every fitted value within 2.1% of the record; adding the differently defined 1595 figure raises the rate to 0.045 yr−1. Set against the base model this is informative: the fitted net expansion of the degree-holding elite is $g_E - \delta = 2.93 \times 10^{-2}$ yr−1, so the two elite layers of Ming society — one recruited by examination, one transmitted by heredity — expanded at nearly the same rate, an order of magnitude faster than the general population. This is reported as an observation about two series, not a claim about a common mechanism, since the rates come from unrelated data.

The landlord compartment L is retained as a structural device but cannot carry empirical content. There is no usable national series of Ming landholding: the quantitative literature consists of local case studies, such as the analysis of a single yulin tuce (fish-scale register) for one county, and no continuous national series of landlord or tenancy counts was found. The compartment L is therefore modelled as a buffer through which a household passes en route from the commoner stock to degree-holding status, and it is tested whether its parameters matter: scanning the mean residence time from 1 to 50 years changes the endpoint commoner share by at most 0.005 percentage points and the fit error by at most 0.07 percentage points. L is dynamically inert, since it only delays a channel far narrower than the stock it feeds, and is presented as a scoping device rather than a modelled quantity.

The examination/heredity split (E_{k}, E_{s}, with privilege feeding the latter at rate σ) appears in Fig. 1b for

completeness of the taxonomy but is neither solved nor calibrated: nothing in the data identifies σ, and no value is assigned to it.

5. Data and Processing

Table 2 summarises the data. All series are generated by scripts from raw sources cached in the project repository, with provenance, checksums and known defects recorded in a data-provenance file; no number in this paper is transcribed by hand.

Table 2 Data series used in the study.

Series	Coverage and size	Source	Notes
Jinshi intake	88 examinations; 24,588	[17]	267 / 6,474 / 17,847 by class; the printed total of the 1514 examination corrected from 376 to 396 (see text)
Origin composition	57 examinations; 15,528	[2]	period means (sample-weighted) 77.95 / 54.54 / 47.34 / 42.38%
Household census	1393; 10,652,870	[15, 16]	initial condition
Elite survival	Ming	[3]	~87% of lineages fail in three generations; gives δ
Licentiate stock	early–late Ming; ≈30,000 → ≈500,000	[19]	external check on $E(t)$; the initial figure is an estimate from the statutory school quotas, not a contemporary count; the terminal figure is Gu Yanwu's late-Ming estimate [19]
Imperial clan	1529–1595; 8,203 → 157,000	[15]	five-class extension; register counts 1529–1565 used for g_R
Cohort check	1550; 320 jinshi	[20]	Section 6.3

The origin series is the one that matters most, and its construction deserves a note. The coding distinguishes three strata: families with no degree and no office in three generations (class A), families holding a licentiate degree but no office (class B), and families with an office-holder (classes C and D) [2]. Because a family holding a degree but no office is not commoner under the model's definition of E, the observable is the class-A share. The broader class-A-plus-B share is reported — the figure usually quoted as "commoner origin" — as a robustness variant in Section 6.4, where the choice of convention is shown not to be neutral.

The intake table as published reports 24,568 graduates whereas its class columns sum to 24,588; the entire 20-name difference arises at the 1514 examination, whose printed total (376) disagrees with its printed breakdown (3 + 135 + 258 = 396). The printed total is a typographical error of the source: the examination-by-examination census of Gong and Qiu gives 396 for that examination [21], and the source's own subtotal for the five examinations of the Zhengde reign (1,773) is consistent only with 396. The 1514 entry is therefore set to 396 and the series total is 24,588. Against the two most recent authoritative counts, 24,586 by Guo Peigui [3] and 24,595 by Gong and Qiu [21], the series differs by two and by seven names respectively; the remaining discrepancies are confined to the 1388 and 1544 examinations (95 against 97 and 312 against 317), where the published counts themselves differ over candidates struck from or restored to the rolls after the palace examination. These are documented disagreements among the sources rather than defects of the series, and the entries are retained as published.

Five of the 57 examinations likewise have internal column sums that disagree with their own usable-sample figure (for example 335 against 326); two independent decodings of the source yield the same figures, so these are defects of the printed source, retained with an explicit flag rather than silently repaired.

6. Results and Discussion

6.1. Parameter estimates and goodness of fit

Fig. 2 plots the fitted trajectory against all 57 observations.

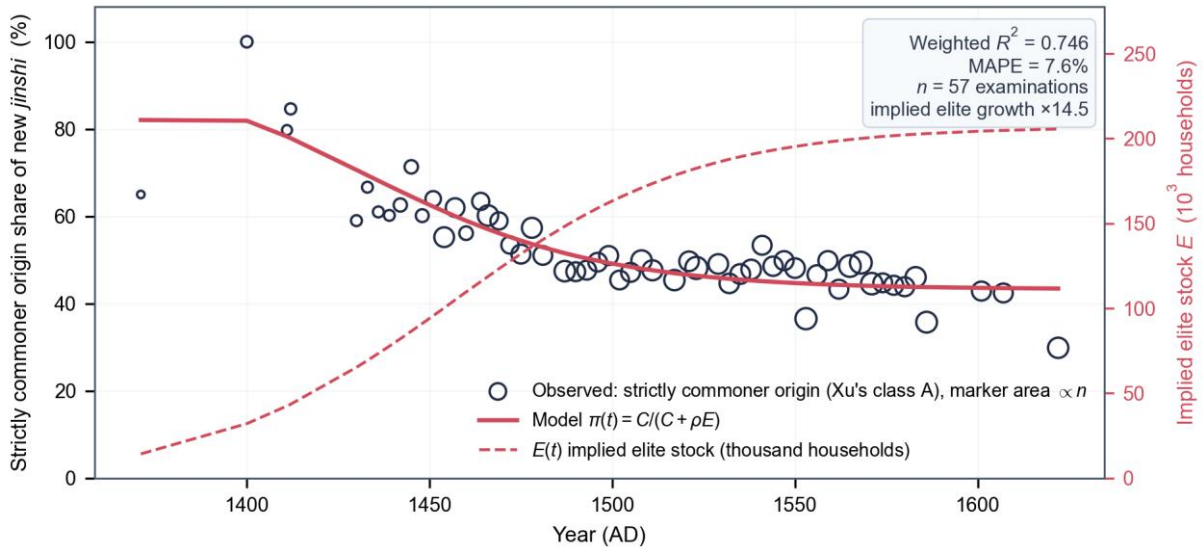


Fig. 2 Fitted trajectory against observation. Open circles: the 57 examination-level observations of the strictly commoner share (marker area proportional to the usable sample size); solid curve: the model $\pi(t)$ from equations (1)–(3). The dashed curve (right axis) is the implied elite household stock $E(t)$ at the fitted point; its level is set by ρ and only its growth factor is identified (Sections 6.1–6.2).

The weighted fit gives $R^2 = 0.746$, a root-mean-square error (RMSE) of 0.0549, a mean absolute error (MAE) of 0.0398 and a mean absolute percentage error of 7.6%. Estimated parameters are $g_C = 9.65 \times 10^{-2}$ yr−1, $\kappa_C = 2.420$, $g_E = 6.33 \times 10^{-2}$ yr−1 (net of decay, $g_E - \delta = 2.93 \times 10^{-2}$), $\kappa_E = 6.822$ and $\rho = 163.7$, with $e_0 = 0.219$; in physical units the commoner capacity is $K_C = \kappa_C C_0 \approx 2.6 \times 10^7$ households and, at the fitted point, the elite capacity $K_E = \kappa_E C_0/\rho \approx 4.4 \times 10^5$ households. These values are one point on a ridge. Re-optimising all other parameters with ρ fixed at values from 30 to 600 changes the weighted R^2 only between 0.746 and 0.747, while g_E stays within 0.062–0.066 yr−1, κ_E within 5.5–8.4 and the growth factor of E over the period within 13.7–20.3; e_0 and the absolute level of E scale inversely with ρ. The reported ρ and the household values derived from it are therefore a scale convention, and only the quantities that are stable along the ridge are interpreted below.

The period means show where the model succeeds and where it does not. For the four conventional periods the observed values are 77.95%, 54.54%, 47.34% and 42.38%; the fitted values are 75.94%, 55.33%, 45.55% and 43.81%. No period deviates by more than 2.1 percentage points, and the four-period profile is reproduced without any period-specific parameter. The residual errors are concentrated at the earliest small examination (1371, $n = 20$), where sample noise is largest and the model's imposed smoothness cannot follow the data.

6.2. Specification choice, and an external check on the implied elite stock

The fitted model contains a capacity parameter K_E on the elite stock, and the first question a reader should ask is whether the data require it. Comparing (2) against the unconstrained alternative $\dot{E} = g_E E$ on the same observable, the constrained specification is preferred on fit quality: R^2 falls from 0.746 to 0.609 and the mean absolute percentage error rises from 7.6% to 9.1% when the ceiling is removed. An elite stock that can accumulate without limit drifts too far from the data over the second half of the period. It is noted explicitly that the external series discussed next does not by itself separate the two specifications — an unconstrained elite can also be tuned to the right growth factor — so the capacity ceiling rests on fit quality, supported by the institutional argument that the supply of degrees and offices was bounded by the quota system [18].

The second question is whether the latent quantity the model infers — the elite stock itself — is consistent with anything known independently. Because the level of E is not identified by π, the meaningful comparison is the growth factor. The fitted model requires the elite stock to rise 14.5-fold between 1371 and 1622 at the fitted point, and between 13.7- and 20.3-fold along the ρ ridge. The independently documented licentiate population rose from roughly 30,000 to 500,000 over the same span, a 16.7-fold increase — the terminal figure being Gu Yanwu's late-Ming estimate [19] and the initial figure an estimate from the statutory school quotas rather than a contemporary count. Because that series never entered the objective function, the agreement is evidential rather than tautological.

Fig. 2 plots this trajectory on the secondary axis. At the fitted point the implied elite share of the commoner stock rises from 0.13% to about 0.8% — a small fraction of the population, which is precisely why a model built on population shares would have been numerically degenerate and historically unverifiable, and why the state variables are household stocks.

6.3. Individual-level cross-validation on a single cohort

The aggregate series can be checked against a single-cohort micro dataset. The 1550 examination (jiajing 29) enrolled 320 jinshi — 3 in the first class, 95 in the second and 222 in the third — and the full roster is recoverable from the China Biographical Database [20]. Independently, Xu Hong's tabulation reports the same 320 usable observations for that examination with a strictly commoner share of 48.13% [2]. The concurrence of the cohort size across two independent sources is itself a useful consistency check on the aggregate series, used here as a validation point rather than as a fitting target. The model predicts 44.7% for 1550, which lies between the register-based figure of 48.13% and the range obtained by applying the published three-generation coding rules to the kin records of the biographical database: 38.6% over all 334 retrieved records, and 40.3% once the 14 records lacking three-generation kin data are excluded from the denominator; all but one of those 14 are the duplicate or variant-character entries that inflate the retrieval beyond the 320-name roster, so the second denominator is, to within one record, the roster itself.

That the two micro codings differ by about 10 percentage points is a property of source scope rather than of either computation being wrong: a biographical database aggregates kin and office information from gazetteers, epitaphs and biographies well beyond the examination register, so an ancestor who held an office is more likely to be recorded — and the household therefore more likely to be classified as elite — than in a register that reports only what the candidate declared. The comparison and its direction are reported rather than the two figures reconciled.

6.4. The definition of "commoner" is not neutral

The same estimation procedure applied to the broader class-A-plus-B share — the figure usually quoted as "commoner origin", which counts families holding a licentiate degree but no office as commoner — gives a visibly better fit: $R^2 = 0.808$ against 0.746, and a mean absolute percentage error of 6.1% against 7.6%. That is not surprising, since the broader series is smoother: the class-B component is negligible before 1550 and grows afterwards, and it is precisely that growth which flattens the broader series in the last century of the dynasty.

The consequence is worth stating plainly. Under the strict convention the commoner share declines continuously and the elite stock must be capacity-limited; under the broader convention the late series plateaus and the elite stock grows by 20-fold rather than 14-fold. Both are internally consistent; they differ in what they call commoner. Since the model's state variable is defined as the set of households holding a degree or an office, only the strict convention aligns the observable with the state. It is therefore adopted as the main specification, with the broader one reported as a robustness variant, but the discrepancy is a reminder that conclusions about whether an elite is saturating depend on where the

boundary of "elite" is drawn.

6.5. Limitations

Three limitations should be stated plainly. First, the elite stock's expansion rate g_E is fitted rather than derived, and the value obtained (net of decay, 2.9×10^{-2} yr−1) exceeds the rate implied by the licentiate series alone (about 1.1–1.4×10^{-2} yr−1). The model channels into E only the jinshi intake and absorbs all other routes into g_E; since the licentiate channel is far wider than the jinshi channel, g_E is best read as a composite of the degree-acquisition process, not as a demographic rate. A specification that adds the licentiate intake explicitly would likely reconcile the two figures, and this is flagged as the natural next step.

The same identification limit applies to ρ, which is set only jointly with the level of E (Section 6.1), and to g_C: the latter is estimated subject to an upper bound of 0.10 yr−1 rather than left free, because π constrains only the ratio $\rho E/C$. C is read as the stock of examination-eligible commoner families, whose growth reflects the expansion of participation and not merely demographic growth; a specification that pins g_C to the demographic rate alone cannot reproduce the observed decline and the licentiate-consistent elite trajectory at the same time, and this is flagged as the model's principal structural limitation. Second, the model is national and ignores regional structure, including the northern/southern register split (Assumption 6) and the province-level quotas that governed provincial examinations [4,18]. Third, the observable is measured among jinshi only — the narrowest and most selective stratum. Whether the same dynamics governed the much larger licentiate population cannot be tested with the present data.

7. Sensitivity Analysis

Fig. 3 reports both a one-at-a-time scan and a global analysis. Four quantities were varied: the degree-quota multiplier q (which scales the elite expansion term g_E), the jinshi-quota multiplier q_m (which scales $m(t)$), the elite decay rate δ, and the entry barrier ρ. Global sensitivity used Latin hypercube sampling with $N = 512$ draws over $q \in [0.6,1.4]$, $\delta \in [0.018,0.055]$ and $\rho \in [0.5\rho^*, 2.0\rho^*]$, with partial rank correlation coefficients computed on the endpoint commoner share, $\pi(1622)$.

Partial rank correlation coefficients were computed in the standard way: each input and the output are rank-transformed, each is regressed on all remaining inputs, and the Pearson correlation of the two residual sets is taken. Ordinary rank correlations, computed for comparison, give the same ranking but overstate the barrier and understate decay.

The ranking is unambiguous, and it is dominated by the two quantities that govern how long elite status lasts and how wide entry into it is. Elite decay is the strongest lever (δ: PRCC +0.970): the faster elite households fall back into the commoner pool, the higher the commoner share remains. The degree quota is a close second and acts the other way (q: PRCC −0.954): widening entry into the degree-holding class enlarges the elite and lowers the commoner share of new jinshi. The entry barrier facing commoner households is a distant third (ρ: PRCC −0.087, $p = 0.049$), an order of magnitude weaker than the other two and only marginally distinguishable from zero in this sample, a consequence of the fact that π constrains the ratio $\rho E/C$ rather than ρ alone.

In the one-at-a-time scan, cutting the degree quota to 60% of baseline raises $\pi(1622)$ by 35.3 percentage points — the system moves towards an almost entirely commoner-origin elite — while raising it above baseline lowers the share by 6.9 percentage points at 140%. The response is strongly asymmetric, because the capacity ceiling binds: once the elite stock is saturated, further degrees cannot concentrate it much more.

The jinshi quota, by contrast, is nearly inert. Doubling q_m moves the endpoint share by only 0.21 percentage points. The response is monotone but trivial, because the annual jinshi intake — on the order of 100 households per year — is an order of magnitude smaller than the elite stock it feeds. This is a substantive finding rather than a computational artefact: the openness of the Ming elite was governed by the width of the entry into the degree-holding class and by the durability of elite status, not by the size of the metropolitan quota. Debates about examination fairness that focus on the number of jinshi awarded are therefore addressing a channel that the dynamics barely transmit.

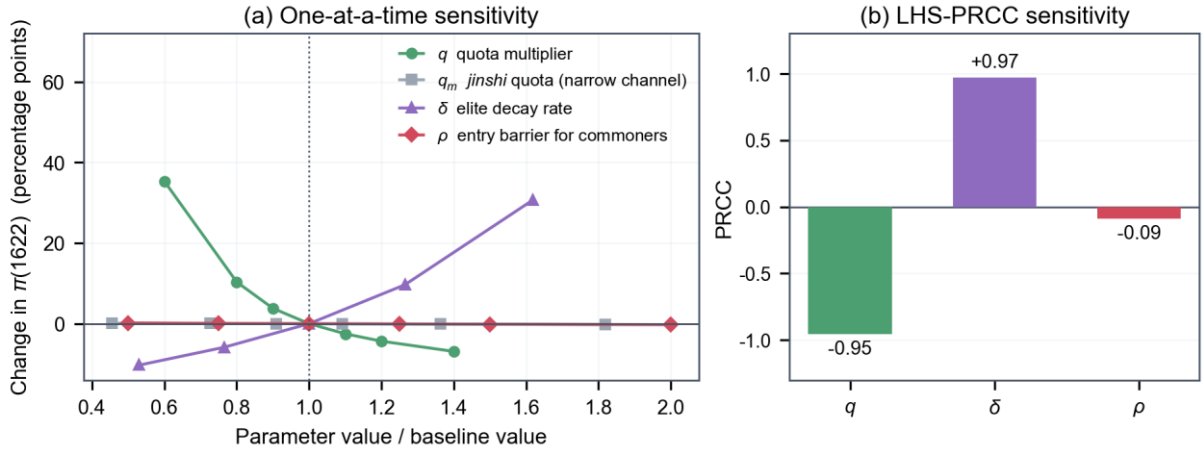


Fig. 3 Sensitivity of the endpoint commoner share $\pi(1622)$. (a) One-at-a-time scans, abscissa normalised by the baseline value of each parameter. (b) Partial rank correlation coefficients from Latin hypercube sampling ($N = 512$).

8. Conclusion

A two-compartment model of Ming social mobility has been built and calibrated, in which the imperial examination acts as a gate with a single behavioural parameter. Calibrated on 57 examinations and validated against an independent series that did not enter the fit, the model reproduces the observed decline in the strictly commoner share of jinshi with $R^2 = 0.746$ and a mean absolute percentage error of 7.6%, and it does so with parameters that are either pinned to documented historical quantities or estimated from the data itself.

Three conclusions follow. First, the observed mobility trajectory requires a capacity constraint on the elite stock; an unconstrained accumulation model fits markedly worse, and the capacity ceiling is the institutional analogue of the quota system that bounded the supply of degrees and offices. Second, the durability of elite status and the width of entry into the degree-holding class are the quantities that govern openness; the entry barrier facing commoner households is not separately identified by the data, and the metropolitan quota is almost inert — a finding with direct implications for how examination reform should be evaluated. Third, the model's latent elite trajectory, inferred only from the commoner-share series, agrees in growth factor (14- to 20-fold against 16.7-fold) with the independently documented rise of the licentiate population, although its absolute level is not identified.

The model is deliberately minimal, and its limitations are correspondingly clear. The most valuable extension is to model the licentiate gate explicitly, which would allow the composite expansion rate to be decomposed and the point of elite saturation to be located rather than fitted.

9. References

[1] Ho, P. T. (1962). The ladder of success in imperial China: Aspects of social mobility, 1368–1911. Columbia University Press. https://cup.columbia.edu/book/the-ladder-of-success-in-imperial-china/9780231051279

[2] Xu, H. (2017). Mingdai xiangshang shehui liudong zai tan [Ming upward social mobility revisited]. Journal of History and Anthropology, 15(1). https://schina-publication.hkust.edu.hk/sites/default/files/pubs/Jour-15.1.A02.pdf

[3] Guo, P. G. (2021). Mingdai de keju gongming qunti ji qi yingxiang yu shehui liudong [Degree-holding groups and social mobility in the Ming]. National Social Science Fund project report. http://www.nopss.gov.cn/n1/2021/0806/c373410-32183931.html

[4] Guo, P. G. (2007). Mingdai juren shuliang ji jinshi pingjun zhongshi nianling kaoshu [The number of Ming juren and the average age of jinshi]. Conference paper. https://ming.cnhis.ncnu.edu.tw/wp-content/uploads/2007/10/Pp12.pdf

[5] Chen, T., Kung, J. K.-S., & Ma, C. C. (2020). Long live Keju! The persistent effects of China's civil examination system. The Economic Journal, 130(631), 2030–2064. https://doi.org/10.1093/ej/ueaa060

[6] Bai, Y., & Jia, R. X. (2016). Elite recruitment and political stability: The impact of the abolition of China's civil service exam. Econometrica, 84(2), 677–733. https://doi.org/10.3982/ECTA13446

[7] Prais, S. J. (1955). Measuring social mobility. Journal of the Royal Statistical Society, Series A, 118, 56–66. https://doi.org/10.2307/2342510

[8] Bartholomew, D. J. (1967). Stochastic models for social processes. John Wiley & Sons.

[9] Shen, D. M. (2006). Ye tan Mingdai qianqi keju shehui de liudong lü: dui He Bingdi yanjiu jielun de sikao [On the mobility rate of early-Ming examination society: reflections on Ho Ping-ti's conclusions]. Shehui Kexue Luntan (Xueshu Pinglun Juan), (9). https://www.aisixiang.com/data/118091.html

[10] Motesharrei, S., Rivas, J., & Kalnay, E. (2014). Human and nature dynamics (HANDY): Modeling inequality and use of resources in the collapse or sustainability of societies. Ecological Economics, 101, 90–102. https://doi.org/10.1016/j.ecolecon.2014.02.014

[11] Turchin, P. (2005). Dynamical feedbacks between population growth and sociopolitical instability in agrarian states. Structure and Dynamics, 1(1).

[12] Alexander, M. A. (2016). Application of mathematical models to English secular cycles. Cliodynamics, 7(1).

[13] Turchin, P., & Nefedov, S. A. (2009). Secular cycles. Princeton University Press.

[14] Kermack, W. O., & McKendrick, A. G. (1927). A contribution to the mathematical theory of epidemics. Proceedings of the Royal Society A, 115, 700–721. https://doi.org/10.1098/rspa.1927.0118

[15] Cao, S. J. (2000). Zhongguo renkou shi, di si juan: Ming shiqi [A history of Chinese population, vol. 4: the Ming period]. Fudan University Press. https://book.douban.com/subject/1584125/

[16] Zhang, T. Y., et al. (1974). Mingshi [History of the Ming]. Zhonghua Book Company. Xuanju zhi, juan 69–71 (schools and examinations); Shihuo zhi, juan 77–82 (household and population registers); together with the imperial genealogy registers, for the household census of 1393 and the recorded clan population.

[17] Zhu, B. J., & Xie, P. L. (1980). Ming-Qing jinshi timing beilu suoyin [Index to the Ming–Qing jinshi steles]. Shanghai Guji Press. https://book.douban.com/subject/1020533/

[18] Da Ming Huidian. (1587). [Collected statutes of the Ming] (Wanli revised edition, 228 juan; statutes on schools and examinations).

[19] Gu, Y. W. (1983). Shengyuan lun [On licentiates]. In H. Chenzhi (Ed.), Gu Tinglin shiwen ji [Collected prose and verse of Gu Yanwu]. Zhonghua Book Company. https://book.douban.com/subject/2243263/

[20] China Biographical Database Project. (2024, February 8). Harvard University, Academia Sinica, & Peking University. https://projects.iq.harvard.edu/cbdb

[21] Gong, Y. M., & Qiu, J. C. (2006). Mingdai dengke jinshi zongshu kao [A study of the total number of Ming jinshi]. Journal of Zhejiang University (Humanities and Social Sciences), 36(3), 69–78. https://www.zjujournals.com/soc/CN/article/downloadArticleFile.do?attachType=PDF&id=8874