The Evolutionary Persistence of Male Androphilia: A Causal-Portfolio Theory of Developmental Paths and Selection Accounting

Tyler M. Moore

Department of Psychiatry, Perelman School of Medicine, University of Pennsylvania; Lifespan Brain Institute, Penn Medicine and Children's Hospital of Philadelphia

Manuscript · October 2026 tymoore@pennmedicine.upenn.edu ORCID 0000-0002-1384-0151

DOI 10.5281/zenodo.23789873

Recommended citation: Moore, T. M. (2026). *The evolutionary persistence of male androphilia: A causal-portfolio theory of developmental paths and selection accounting*. Manuscript retrieved from https://mooremetrics.com/homosexuality. https://doi.org/10.5281/zenodo.23789873

Abstract

The conventional evolutionary paradox of male homosexuality is mis-specified because it combines distinct developmental and population-genetic accounting problems. Etiological heterogeneity does not remove the selective burden carried by inherited liability; it determines which burden belongs to which causal component. This article advances a causal-portfolio theory in which enduring male androphilia is defined by repeated male-directed and female-directed attraction rather than by identity, partnership, or behavior. Four definitions fix the explanandum, five causal levels, and one reporting scale. Four empirical propositions are placed at risk: androphilia is a convergent phenotype produced by at least two mechanism-distinct prenatal paths; its inherited architecture is distributed and shared across sexes; prenatal organization, not a scalar androgen deficit, is its principal causal stage; and its maintenance is component-specific, with inherited liability compensated in aggregate across full-genome partitions covering at least 60% of common-SNP heritability, the verdict confined to that covered majority, and recurrent prenatal exposures costly to remove. A path qualifies only if its standardized effect on latent attraction is at least .10 with a lower 90% bound above .05 in an independent hold-out, it survives an assumed residual mediator-outcome correlation of .20, and a design independent of sequential ignorability supports it. The core is rejected only when fewer than two paths qualify and a route-free comparator is equivalent to the best path model in a pooled hold-out; the operating characteristics of that decision are reported. Six critical studies separate failure of the core from failure of a subarchitecture, path, or assay.

Keywords: Male androphilia; Sexual orientation; Causal inference; Evolutionary maintenance; Falsifiability

Contents

Introduction

The Causal-Portfolio Theory

Formal Causal Architecture

The Decision Rule and Its Operating Characteristics

Competing Models and Discriminating Predictions

A Worked Example: Same Data, Different Identification Assumptions

Empirical Constraints on the Causal Graph

Evolutionary Accounting

Critical Tests: How the Theory Can Fail

Evolutionary Synthesis

Conclusion

Declarations

References

Supplementary Material

S1 Shared Genetic Architecture and Sex-Differential Expression

S2 Twin, Family, and Registry Estimates in Detail

S3 The Xq28 Linkage Record and Its Ascertainment Dependence

S4 Non-European Ancestry and Maternal X-Regulatory Variation

S5 Population-Genetic Models and Evidence for Carrier Advantage

S6 Maternal-Line Fecundity: Supporting and Contradictory Evidence

S7 Kin-Directed Investment: Positive and Negative Evidence

S8 Childhood Gender Nonconformity: The Bivariate Twin Evidence

S9 The Fraternal Birth-Order Association: Quantification and Alternatives

S10 The Postmortem and Neuroimaging Record

S11 Causal Paths, Prenatal Organization, and Model Comparison

S12 Comparative Mate-Choice Theory and Speculative Social Extensions

S13 Formal Causal-Portfolio Accounting and Full Critical-Study Protocols

Supplementary References

Introduction

Male homosexuality is evolutionarily consequential because enduring androphilia has often reduced direct reproduction, shows familial and genetic influence, and nevertheless recurs across populations and historical periods (Apostolou, 2022; Bailey & Pillard, 1991; Bailey et al., 2000; Kendler et al., 2000). The burden is real. What is wrong is the assumption that the phenotype, its causes, and its persistence can all be represented by one inherited character at one equilibrium. That formulation demands one answer where the biology presents several different accounting problems. Throughout this article, androphilia names the repeated attraction construct in Definition 1. When a cited study measured homosexuality, same-sex behavior, identity, partnership, or legal union, its own phenotype term is retained and is not treated as interchangeable with androphilia.

Four questions must be separated. First, which inherited variants or polygenic components raise liability, and why do they remain polymorphic? Second, which maternal or fetal developmental systems can generate the phenotype, and why are the relevant vulnerabilities not eliminated? Third, why do noninherited pregnancy-dependent events recur in successive generations? Fourth, how does the joint distribution of inherited liability, maternal physiology, fetal exposure, and developmental regulation generate the observed prevalence in a particular cohort? Heterogeneity divides these burdens; it does not erase any of them.

The correct opening claim is therefore exact: the conventional paradox is mis-specified because it combines several distinct evolutionary and developmental accounting problems. An inherited component that substantially raises the probability of reduced reproduction still requires a population-genetic explanation. A recurrent maternal immune exposure does not. It instead requires an incidence model and an explanation for the persistence of the causal immune specificity, placental transfer process, fetal target, or downstream sensitivity. A theory that merely lists both processes has not solved either one.

The unit of inherited accounting is the variant or molecular component and its average fitness effect across every body in which it occurs. Female phenotypes remain inside that ledger because men and women share much of the same inherited architecture. The largest behavioral GWAS estimated a substantial cross-sex genetic correlation, approximately .63, even though expression differs by sex (Felesina, 2026; Ganna et al., 2019). The developmental explanandum can be male-specific without the evolutionary bookkeeping becoming male-only.

The theory also makes no group-selection claim. Coalition effects, pair bonding, kin investment, and social organization matter only when they change the survival or reproduction of carriers or genetic relatives and can be attributed back to inherited variation. Comparative evidence can show that same-sex behavior is common, plastic, or socially functional; it cannot by itself explain the persistence of stable human male-directed attraction.

The Causal-Portfolio Theory

Stated plainly, the theory holds that enduring male androphilia is a single endpoint reached by more than one developmental path. Inherited liability, maternal physiology, pregnancy history, and developmental chance act through prenatal endocrine, placental, immune, and regulatory mechanisms; those mechanisms organize how male and female cues are valued; and the resulting attraction profile is what behavior, identity, and partnership imperfectly express. Because the paths differ in kind, their persistence must be explained in different currencies. An inherited component is answerable to allele-level selection summed over every carrier of either sex; a recurrent pregnancy-dependent event is answerable to an incidence model and to the cost of removing the vulnerability it exploits. The theory is a portfolio in exactly that sense: the phenotype is produced by several holdings, each kept on its own ledger, and the maintenance verdict is rendered on the measurable aggregate of the inherited holdings rather than on any one favorable component. The four questions above are divided accordingly: the first falls to the inherited arm of P4; the second and third to its recurrent arm, which pairs a constraint test with an incidence model; and the fourth to the joint-distribution definition D4, while P1 through P3 state what the developmental architecture must be for that division of burdens to be the right one.

The theory is stated as four definitions and four empirical propositions. The distinction is deliberate and is not a matter of presentation. The definitions fix what is being explained, what counts as a route, and on which scale quantities are reported; they can be judged useful or useless, but they cannot be false. The empirical propositions can fail, and the critical studies below state how. Table 1 collects every prespecified quantity, threshold, and margin in one place; the text refers to Table 1 rather than restating the definitions at each use.

Table 1 *Prespecified Quantities, Thresholds, and Decision Margins*

Quantity	Prespecified value	Scale or basis
Explanandum region *R*	$\bar{A}_M \geq +1.80$ and $\bar{A}_F \leq -0.75$	Male-population standardized longitudinal means
Measurement rule *m*	Person-level means over $T \geq 4$ occasions spanning ≥ 5 years, then region *R*	Declared deterministic function of the latent summaries
Planning prevalence π	.025 (sensitivity .017 to .036)	Implied by *R* under the planning joint distribution
Planning attraction correlation	−.55 (sensitivity −.45 to −.65)	Bivariate-normal planning approximation for $(\bar{A}_M, \bar{A}_F)$
Materiality threshold	Standardized effect ≥ .10	Latent attraction
Lower-bound gate	Lower 90% bound > .05 in an independent probability-weighted hold-out	Latent attraction
Robustness value ρ^*	≥ .20	Assumed absolute residual mediator-outcome correlation
Independent identification	At least one design not resting on sequential ignorability	Within-family or within-mother contrast, instrument, targeted perturbation, or negative control
Practically null	< .03	Latent attraction
Indeterminate zone	.03 to .10	Latent attraction
Primary causal equivalence margin	.03	Latent attraction; two one-sided tests at 90% confidence
Derived prevalence margin	.0015 under the planning assumptions; recomputed from the fitted joint distribution	Image of the .03 causal margin under *m*; $d\pi/d\delta = .048$
Secondary predictive margins	$\Delta R^2 < .0025$; ΔBrier < .002; Δ log-loss < .005	Held-out prediction; secondary, never decisive alone
Rejection rule	Fewer than two qualifying paths **and** equivalence on the causal margin plus at least three of four secondary margins	Pooled probability-weighted hold-out
Stratification	Four reference-projected genetic-similarity strata; self-reported ancestry retained as a covariate and for invariance testing	Heterogeneity diagnostics, not gates (National Academies of Sciences, Engineering, and Medicine, 2023)
Consequential bound, Critical Study 3	.10 with lower 90% bound above .05	Identical to the materiality gate
Coverage criterion *C*	≥ .60 with lower 95% bound ≥ .50	Fraction of out-of-sample common-SNP heritability in selection-valid partitions
Planning h^2_{SNP}	.15 (planning range .08 to .25); estimated in the discovery sample	Repeated attraction; the coverage denominator is estimated, not assumed
Compensation criterion	Lower bound on $G_C/K_C \geq 1.0$, or a separately fitted nonnegative equilibrium	Covered partitions only; no imputation to $1 - C$

Quantity	Prespecified value	Scale or basis
Planning design, Critical Study 1	10,000 analyzed men per stratum (40,000 total); 30% hold-out	Smallest grid point meeting the operating characteristics below
Target operating characteristics	≥ .90 correct retention; ≥ .80 correct rejection; familywise false-decision probability ≤ .05	End-to-end Monte Carlo simulation of the complete decision

Note. All thresholds are fixed before data collection; where the text repeats a value, this table is authoritative. Standardized effects refer to mediated or identified edge-intervention effects on latent attraction. The robustness value is the smallest assumed absolute residual mediator-outcome correlation at which a path's lower 90% bound falls to .05.

Definitions and Commitments

D1. The explanandum is enduring male-directed attraction, not a broad social category. Attraction to men and attraction to women are measured separately and repeatedly. Fantasy, behavior, identity, partnership, and legal union are downstream or imperfect indicators. Predominant or exclusive male androphilia is the prespecified region R of the joint attraction distribution given in Table 1 and shown in Figure 2, not a synonym for any single observed indicator.

D2. Phenotypic distinctness and genetic separability are different claims. Orientation, behavior, identity, and partnership can remain phenotypically noninterchangeable even if one genetic factor contributes to all of them. The separate genetic claim, that cue-specific attraction and general behavioral propensity retain distinguishable inherited components after reliability and cross-sex covariance are modeled, appears as P2 and is tested in Critical Study 2. Failure of P2 does not make the observed phenotypes identical.

D3. Five causal levels are not collapsed. Offspring genotype, maternal genotype and physiology, pregnancy exposure, and stochastic variation are sources. Steroid signaling, placental transport, immune activity, and gene regulation are prenatal mechanisms. Sex discrimination, cue-specific incentive salience, attachment, and social reward are developmental intermediates. Repeated male-directed and female-directed attraction constitutes a distinct latent-attraction level. Fantasy, behavior, identity, partnership, and categorical classification are observed phenotypes or measurement rules. Diffuse polygenicity describes architecture, maternal-line concentration describes a transmission signature, and regulatory state describes a mechanism or modifier. None is automatically a route.

D4. Prevalence is a joint-distribution expectation, and all margins are stated on one scale. Cohort prevalence is the expectation of the outcome probability over inherited liabilities, maternal states, exposures, mechanisms, and their interactions. It is not the arithmetic sum of named route frequencies. This is descriptive standardization, not a theoretical result. Because the explanandum is latent attraction, every primary estimand and every equivalence margin is defined on the latent-attraction scale; quantities defined on the observed classification Y are derived from those through the fixed rule m, as set out below.

Empirical Propositions

P1. Male androphilia is a convergent phenotype produced by more than one mechanism-distinct prenatal path. The claim is stronger than generic multifactorial causation. At least two directed source-to-mechanism-to-intermediate paths must make material, replicable contributions to enduring attraction and must generate different temporal, conditional-independence, external-prediction, or intervention-relevant patterns. Material is numerical and is defined once, in Table 1, as a standardized mediated or identified edge-intervention effect on latent attraction of at least .10; a path qualifies only if it also shows a lower 90% confidence bound above .05 in an independent probability-weighted hold-out, a robustness value of at least .20, and support from at least one design that does not rely on sequential ignorability. Overlapping loadings and interactions are expected; sealed etiological classes are not.

P2. The inherited architecture is distributed, possibly omnigenic, and substantially shared across sexes, but its phenotypic channels need not be identical. No common high-penetrance locus is necessary or sufficient. Many small effects may enter several prenatal mechanisms and may be correlated through linkage disequilibrium, pleiotropy, or genome-wide regulatory coupling. The theory therefore does not require a small set of orthogonal inherited factors, and its evolutionary test partitions the full common-variant signal into mutually exclusive, preregistered bookkeeping bins that can each contain thousands of correlated effects.

P3. Prenatal organization is the principal causal stage, not a one-dimensional deficit in androgen exposure. Major biological effects on enduring attraction are established through prenatal or early perinatal organization of cue valuation and social reward. The proposition does not predict global under-masculinization, one direction of hormone effect in every subgroup, or a single anatomical signature. Steroid concentration, receptor sensitivity, placental metabolism, immune perturbation, cell-specific regulation, and developmental timing can generate different configurations.

P4. Evolutionary maintenance is component-specific and carries an explicit aggregate commitment. Inherited variation may persist through antagonistic pleiotropy, sexual antagonism, overdominance, mutation-selection balance, drift, linkage, or another quantitatively specified process, and different parts of the genome may be maintained by different processes. The inherited commitment is made at the level of the covered aggregate, not of any single favorable component: the compensatory-maintenance arm is tested with selection-valid, mutually exclusive full-genome partitions covering the prespecified majority of common-SNP heritability given in Table 1. Their carrier-weighted aggregate under natural fertility must be nonnegative after expression costs are included. The uncovered fraction is left unclassified: no effect estimated in the covered set is extrapolated to it, and it cannot rescue or condemn the covered portfolio. Pregnancy-dependent and stochastic paths do not require balancing selection on orientation, but the implicated vulnerability must be costly or constrained to remove.

Figure 1 maps the five causal levels and highlights two contrasting exemplar paths against the broader candidate graph.

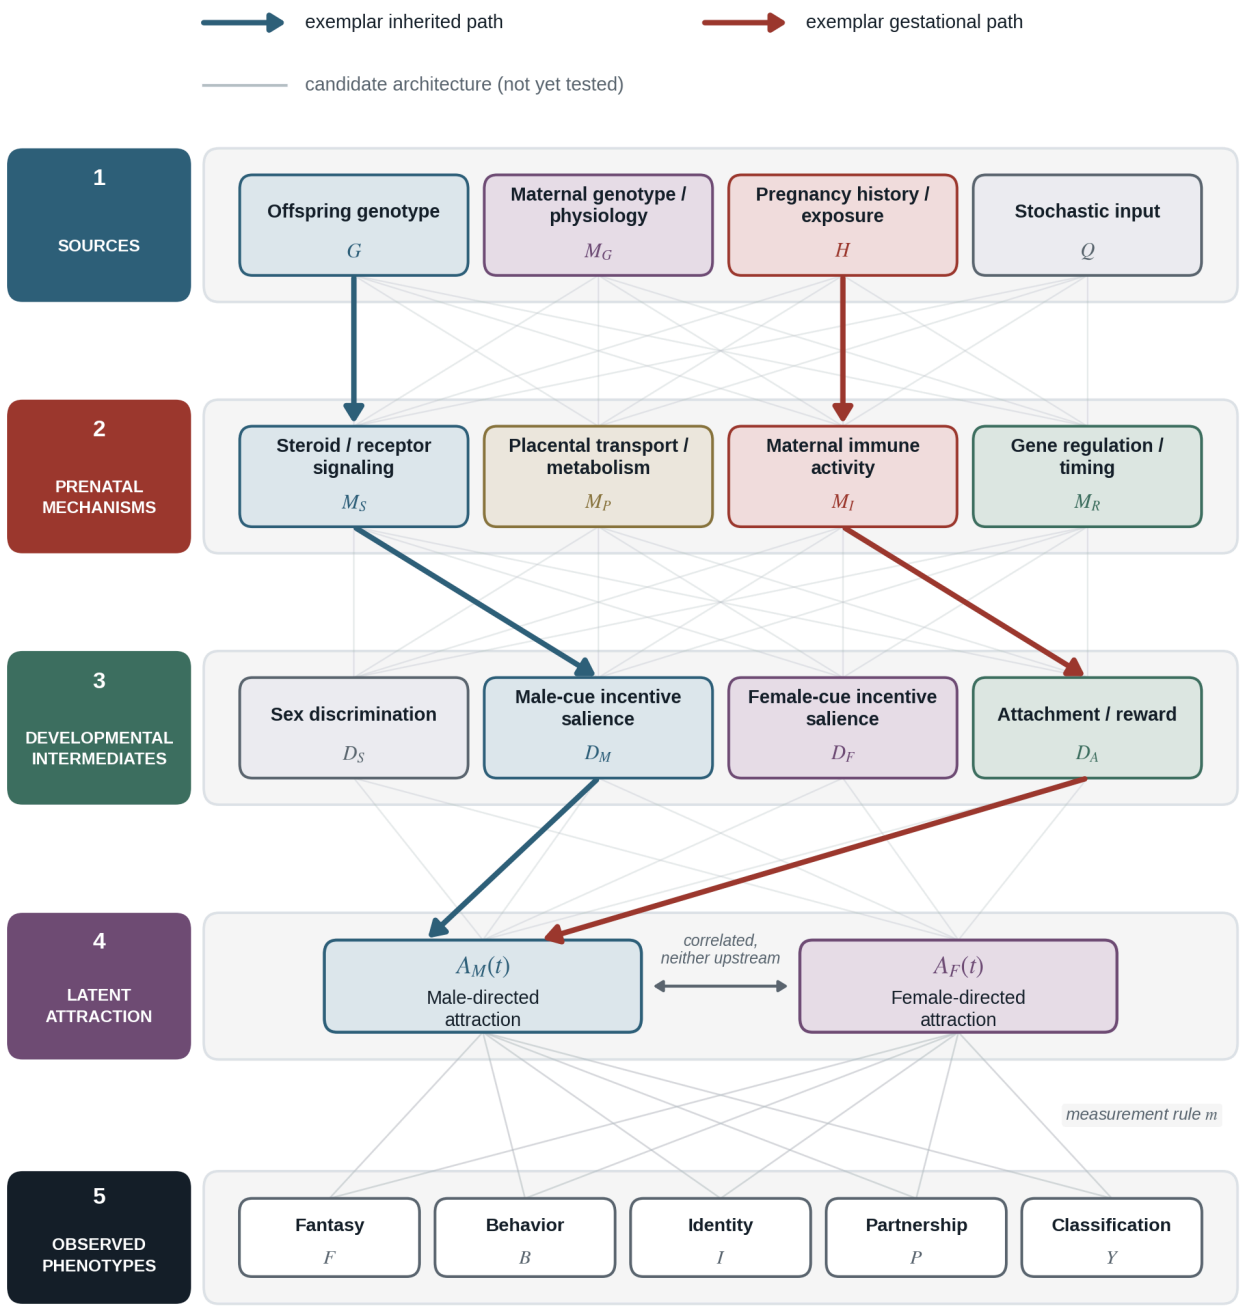


Figure 1 *The Causal-Portfolio Architecture: Five Levels and Two Exemplar Paths*. Sources, prenatal mechanisms, developmental intermediates, latent attraction, and observed phenotypes are separate levels. Grey edges are candidate architecture; highlighted edges are exemplar test paths. Both exemplar paths terminate at the latent attraction level, because a route is defined from a source through mechanisms and intermediates to attraction; the grey edges below that level represent the measurement rule, not causation. A pedigree pattern or genomic architecture is evidence to explain, not a path definition.

Formal Causal Architecture

Sources, Mechanisms, and Paths

Let S_i denote sources for male i: offspring inherited genotype, maternal genotype and physiology, pregnancy history and exposure, and stochastic developmental input. Let M_i denote prenatal mechanisms, D_i developmental intermediates, and $A_i = (A_{Mi}, A_{Fi})$ the latent attraction state. Observed phenotypes O_i include fantasy, behavior, identity, partnership, and a categorical rule Y_i. A minimal acyclic representation is:

$$M_i = f_M(S_i, U_{Mi}); D_i = f_D(M_i, S_i, U_{Di}); A_i = f_A(D_i, M_i, S_i, U_{Ai}).$$

A route is a directed path from one element of S_i through one or more elements of M_i and D_i to A_i. Two routes are causally distinct only if the data support nonredundant temporal order, mediator dependence, conditional-independence restrictions, identified edge interventions, or external predictions. Mediation alone is insufficient because modest unmeasured mediator-outcome confounding can manufacture the same signature. A path therefore requires both confounding-robust prospective mediation and an independent identification source (the two-key rule). Different predictors loading on different factors is insufficient. Conversely, two causal routes may converge onto one smooth attraction distribution, so failure to recover latent classes does not demonstrate one etiology.

This graph resolves the causal-level problem. Diffuse polygenicity can feed steroid signaling, immune susceptibility, placental biology, neural gene regulation, or several simultaneously. Maternal-line concentration can arise from X linkage, maternal genotype, imprinting, interacting autosomal loci, correlated maternal exposure, or ascertainment. Maternal immune activity is a mechanism linked to pregnancy history. Regulatory state may mediate or modify any path. The directed graph, not the label, determines the route.

The Primary Phenotypes and the Classification Rule

For each measurement occasion t, let $A_{Mi}(t)$ denote latent male-directed attraction and $A_{Fi}(t)$ latent female-directed attraction. Each is estimated from multiple indicators of intensity, physiological or subjective response, fantasy, persistence, and partner-specific incentive value. The primary adult quantities are the longitudinal summaries $\bar{A}_{Mi}$ and $\bar{A}_{Fi}$, the person-level means across $T \geq 4$ occasions spanning at least five years, standardized in the male population. They are not a one-dimensional Kinsey score and not an ever-versus-never behavior item.

The measurement rule m and the region R are specified here rather than named. The rule maps the repeated profile to its person-level means and applies the region:

$$Y_i = m(\{A_{Mi}(t), A_{Fi}(t)\} \text{ for } t = 1, \ldots, T) = 1 \text{ if and only if } \bar{A}_{Mi} \geq +1.80 \text{ and } \bar{A}_{Fi} \leq -0.75.$$

Under the planning joint distribution described below, R contains approximately 2.5% of men, consistent with published estimates of predominant or exclusive androphilia in probability samples. The thresholds, the required number and span of occasions, and the standardizing population are fixed in advance and are reported in Table 1. The rule is a declared function of the latent summaries, not a separate construct: R is a corner of one continuous joint distribution, not a class recovered from the data.

Same-sex behavior is downstream. Let L_i index general mating effort and behavioral propensity, V_i opportunity and partner availability, C_i social, legal, or interpersonal constraint, and U_{Bi} unmeasured behavioral inputs. Then behavior is represented as:

$$B_i = g(A_{Mi}, A_{Fi}, L_i, V_i, C_i, U_{Bi}).$$

This removes the circularity in which partner-sex valuation was placed inside the function that defined orientation. Attraction is modeled directly; behavior is an outcome of attraction plus general propensity and context. A one-factor genomic result can therefore reject P2 without implying that attraction, behavior, identity, and partnership are interchangeable measurements.

Figure 2 shows the joint attraction distribution with the region R marked, and places the downstream behavioral and classification indicators.

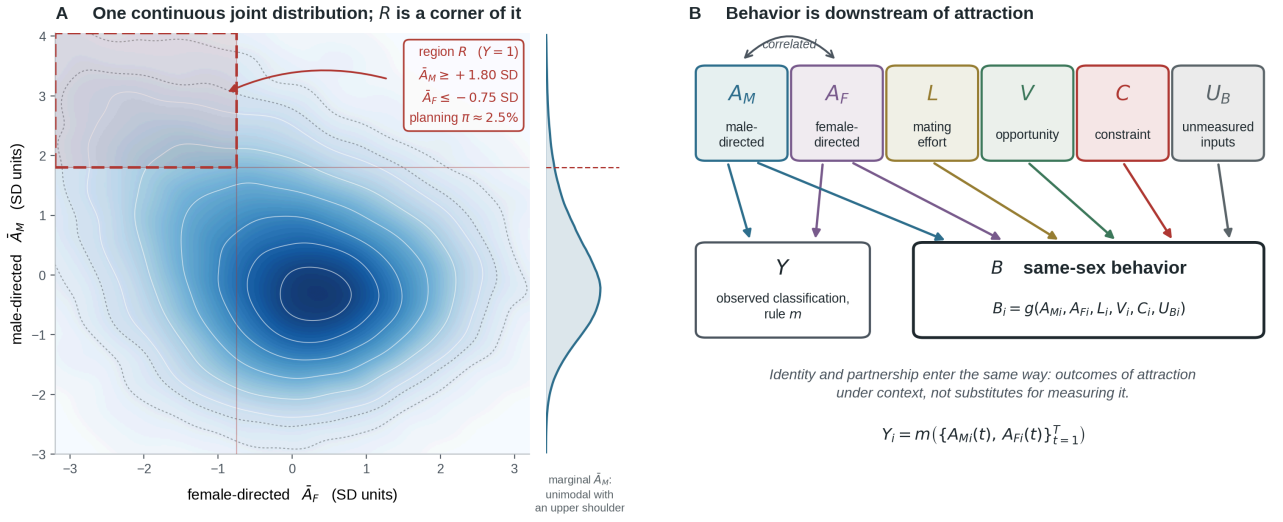


Figure 2 *The Phenotype Model: Repeated Attraction Dimensions and Downstream Behavior*. Panel A shows an illustrative joint distribution of the two repeated attraction summaries as a single continuous surface. The marginal distribution of $\bar{A}_M$ at the right is unimodal with an upper shoulder rather than bimodal, and the prespecified region R, whose thresholds are given in the callout, is a corner of that continuum rather than a separated cluster. This is the point of the panel: one smooth surface is compatible with several causal routes, so failure to recover latent classes is not evidence for one etiology. Panel B places same-sex behavior downstream from attraction, general mating effort, opportunity, and constraint. The genetic dimensionality of these quantities is an empirical question distinct from their phenotypic meaning.

Prevalence, Node Interventions, and Path-Specific Effects

Let Z contain measured sources, mechanisms, intermediates, covariates, and interactions, and let F_t be their joint distribution in cohort t. Population prevalence is:

$$\pi_t = E_{Z \sim F_t}[\Pr(Y = 1 \mid Z)].$$

This expression is descriptive standardization over the cohort joint distribution. It is not a controlled-direct-effect formula, and the presence of intermediates inside Z does not make conditioning on them causally innocuous. The expression permits nonlinearities, overlap, and interaction. It also prevents the contradiction created by calling prevalence a simple sum of route incidences. Inclusion-exclusion is available only when routes are represented as binary sufficient-cause events. With continuous loadings and interaction, a case produced by moderate inherited liability plus moderate immune exposure has no unique natural assignment to either route.

A node intervention is a total effect of changing a source or mechanism, not a path-specific effect. For a prespecified node X and reference x_0, define $\Delta_{\text{node}}(X, t) = \pi_t - \pi_t\{\text{do}(X = x_0)\}$. Under this sign convention, a protective reference intervention that lowers prevalence yields a positive Δ_{node}. This quantity inherits the standard assumptions of structural causal models and target-population transport (Hernán & Robins, 2020; Pearl, 2009).

For a prespecified set of directed edges P, the primary edge-intervention contrast is defined on latent attraction:

$$\text{PSE}_{P,t} = E_t[\bar{A}_M(x \text{ along } P, x_0 \text{ along } \bar{P}) - \bar{A}_M(x_0)].$$

The corresponding contrast on Y is reported as a derived quantity obtained by applying m. The node, edge set, mediator order, target population, and treatment of interaction must be declared in advance. Edge-specific effects are not identified merely because all nodes are measured. A recanting witness or exposure-induced mediator-outcome confounder can block identification precisely when interacting paths share downstream nodes (Avin et al., 2005; Robins & Greenland, 1992; VanderWeele et al., 2014). When edge-identification conditions fail, report node or joint-intervention effects, interventional analogues, or bounds rather than path-attributable cases (Imai et al., 2010; VanderWeele, 2015). Shapley allocation is permissible only as a convention-dependent descriptive decomposition (Shapley, 1953).

The Robustness Gate Is a Curve, Not a Pass or Fail

Every mediated-effect decision includes a prespecified residual-correlation sensitivity analysis. Rather than a single assumed ceiling, each candidate path is summarized by its **robustness value** ρ^*, the smallest assumed absolute residual mediator-outcome correlation at which the path's lower 90% bound falls to .05 (Cinelli & Hazlett, 2020; VanderWeele & Ding, 2017). Under

the linear model used in the worked example, the bias-adjusted standardized indirect effect is approximately $\theta\hat{} - a\kappa\rho$, where a is the source-to-mechanism coefficient and κ is the ratio of the outcome to mediator residual standard deviations. The gate is $\rho^* \geq .20$.

Two consequences must be stated rather than left implicit. First, the materiality threshold and the robustness gate are not independent. With $a = .40$ and $\kappa = 1.02$, a path whose naive indirect effect is exactly .10 has $\rho^* = .12$ (these planning values treat sampling error in the hold-out as negligible); requiring $\rho^* \geq .20$ therefore implies a naive indirect effect of about .132. The effective materiality threshold is a function of the gate, and both are reported. Second, a stress test fixed at an assumed residual correlation of .10, the ceiling adopted in an earlier version of this framework, is too lenient. The justification for .20 is given with the worked example below, and Figure 5A reports the full curve so that readers can apply their own ceiling instead of accepting a single pass-or-fail verdict.

Placing All Margins on One Scale

Materiality and the causal equivalence margin are defined on latent attraction A. Prevalence, Δ_{node}, and the prevalence equivalence margin are defined on $\pi = \Pr(Y = 1)$, the output of the classification rule. These are not interchangeable scales, and a margin that looks strict on one can be permissive on the other. Because m and R are now fixed, the mapping can be computed rather than assumed.

For planning, the joint distribution of $(\bar{A}_{\text{M}}, \bar{A}_{\text{F}})$ is approximated as bivariate normal with correlation −.55. Under that approximation and the rule in Table 1, $\pi = 2.53\%$, and the local derivative of prevalence with respect to a uniform shift δ in $\bar{A}_{\text{M}}$ is $\mathrm{d}\pi/\mathrm{d}\delta = .048$. A standardized effect of .03 on latent attraction therefore corresponds to a prevalence difference of .0015, and the materiality threshold of .10 corresponds to .0048. Across plausible values of the attraction correlation (−.45 to −.65) and threshold pairs yielding π between 1.7% and 3.6%, the derivative ranges from .036 to .058 and the derived margin from .0011 to .0017.

This computation shows why the prevalence margin cannot be fixed independently. A prevalence equivalence margin of .005, which looks strict and was the value fixed in an earlier version of this framework, corresponds to .103 standardized units on latent attraction, approximately the materiality threshold itself and about 3.4 times the causal margin of .03. Two models differing by a full material effect would have satisfied it. The prevalence margin is therefore derived as the image of the .03 causal margin under m, equal to .0015 under the planning assumptions and recomputed from the fitted joint distribution in the analysis. The remaining predictive margins are secondary, are reported on their own conventional scales, and never carry the primary decision.

This also resolves an apparent tension with D1. The explanandum is latent attraction; Y is a declared deterministic function of the latent summaries; quantities on Y are therefore derived

reports of the same underlying causal structure rather than a second, categorical explanandum.

The Decision Rule and Its Operating Characteristics

The causal core is rejected when both of the following hold in the pooled, probability-weighted hold-out: fewer than two paths pass the materiality-and-identification gate, and the route-free comparator is statistically equivalent to the best path model on the primary causal margin and on at least three of the four secondary predictive margins. Equivalence is assessed with two one-sided tests at 90% confidence (Lakens et al., 2018). The route-free comparator is defined as a structural causal model with exactly one or two latent continuous liabilities, no source-specific mediator map, and no identified edge effects; three- to five-liability versions are mandatory sensitivity models. The core is retained only when at least two paths pass the gate and the comparator fails equivalence; every other pattern, such as a single qualifying path without comparator equivalence, is reported as indeterminate rather than as retention.

Three features of this rule depart from the more obvious design, a conjunctive rule that requires equivalence on every margin within every stratum, which an earlier version of this framework used and which cannot fire at any feasible sample size.

First, the primary decision is made in the pooled hold-out rather than separately within every genetic-similarity stratum. Requiring equivalence within every stratum makes the least-powered stratum the binding constraint and gives the perverse result that adding strata makes the theory harder to reject. Stratum-specific estimates are now reported as heterogeneity diagnostics. If stratum-specific point estimates for the primary contrast differ by more than the causal margin, the pooled verdict is withheld and the result is reported as heterogeneous rather than as retention.

Second, equivalence is required on the primary causal margin and a prespecified majority of the secondary margins rather than on every margin conjunctively. Establishing equivalence is harder than failing to reject a difference, because a two one-sided test requires the whole interval to fall inside the margin. Requiring roughly two dozen such conclusions, every margin in each of four strata, guarantees indeterminacy rather than falsification.

Third, the operating characteristics of the rejection decision are reported alongside those of retention. This is the substantive point. Under the conjunctive rule, with 5,000 analyzed men per stratum and a 30% hold-out, the half-width of the 90% interval for the primary contrast is $1.645 \times .026 = .042$, which exceeds the .03 margin, so no dataset of that size can declare equivalence at all. Simulation across a design grid confirms the consequence: the probability of a correct rejection under a route-free data-generating process was below .001 for every design up to 80,000 analyzed men and did not reach .80 until roughly 250,000. Under the pooled rule the same probability is .21 at 20,000, .84 at 40,000, and .98 at 60,000 (Figure 5B). The planning design is therefore 10,000 analyzed men in each of four strata, a total of 40,000, with a 30% hold-out.

The final sample is the smallest value in a preregistered grid for which end-to-end Monte Carlo recovery of the complete decision procedure reaches at least .90 probability of a correct retention under a portfolio generator and at least .80 probability of a correct rejection under a route-free generator, with familywise false-decision probability at or below .05. The simulation must include the case oversample, design weights, hold-outs, stratum diagnostics, missingness, nonnormality, measurement error, interaction, ascertainment, and the mediator-outcome confounding scenarios used in the robustness gate.

Feasibility follows from the prevalence and can now be derived rather than asserted. Enrolling 10,000 men in the region R, a 25% oversample of the analyzed sample, requires screening approximately 395,000 eligible responders at $\pi = 2.5\%$. At a 20% screening response and 50% enrollment, that is roughly 4.0 million invitations; at a 25% response and 60% enrollment, roughly 2.6 million. A design with a smaller oversample of 5,000 requires between 1.3 and 2.0 million. These figures are large, and they are stated rather than hidden.

A note on what is claimed. The theory is falsifiable by prespecified adverse patterns whose operating characteristics are stated above. Critical Studies 1 through 4 and 6 are runnable with current technology at the sample sizes given. Critical Study 5 is not, because no existing resource combines a genome-wide study of repeated attraction with directly genotyped natural-fertility outcomes; it is a long-horizon falsifier. Describing the whole program as decisively falsifiable would overstate what is presently available, and the distinction is maintained throughout.

Competing Models and Discriminating Predictions

The causal-portfolio theory extends, but does not originate, the proposition that similar adult orientations may arise through multiple biological processes. Rahman and Wilson (2003) synthesized genetic, prenatal endocrine, neurodevelopmental, and evolutionary evidence without a path-identification framework. VanderLaan and colleagues later argued explicitly for biodevelopmental subgrouping and for using gender expression and other markers to carve orientation at its joints (VanderLaan et al., 2022). The present theory is distinct only if it earns the additional commitments listed below: routes must be directed paths through a causal graph, phenotype production must be separated from observation, node effects must not be mislabeled as path effects, and each inherited or recurrent component must be paired with its own evolutionary accounting rule.

Table 2 states the closest alternatives and applies one decision logic throughout. The theory is not supported merely because several predictors are significant, several factors fit, or several mechanisms are biologically plausible. It survives only when at least two paths pass the materiality-and-identification gate and the joint pattern rejects the alternatives.

Table 2 *Competing Models and the Observations That Discriminate Among Them*

Model	Causal claim	Discriminating prediction	Decision pattern
Causal-portfolio theory	At least two mechanism-distinct prenatal paths contribute materially; paths may overlap and share downstream nodes.	Each path clears the materiality-and-identification gate in Table 1; inherited and recurrent components require different accounting.	Retained only if at least two paths pass that gate and route-free comparators fail equivalence under the pooled rule.
Route-free continuous liability	Several continuous dimensions influence one developmental outcome, without separable causal paths.	Dimensions improve fit, but no material, confounding-robust, independently identified source-specific path remains.	Favored when one or two liabilities are equivalent to the best path model on the causal margin and a majority of secondary margins.
Biodevelopmental profile heterogeneity	Biomarker-defined developmental profiles or subgroups contribute to similar orientations.	Subgroups or profiles replicate, but evolutionary maintenance and counterfactual prevalence effects need not be specified.	Favored if replicated profiles do not yield additional causal, prevalence, or maintenance predictions.
Broad polygenic architecture plus measurement	One distributed inherited architecture produces variation; apparent routes reflect phenotype definition, ascertainment, or proxy error.	Direct mechanism measures add no material, confounding-robust prediction beyond polygenic liability and reliable attraction measures.	Favored if endocrine, immune, placental, and regulatory indicators produce no qualifying path.
Independent genetic and gestational mechanisms	Genetic and maternal mechanisms both exist but require no integrated portfolio or interaction accounting.	Main effects replicate, but overlap, interaction, and joint-distribution predictions are unnecessary.	Favored if the integrated model provides no additional calibrated prediction or counterfactual utility.
Mutation-selection without compensation	Inherited liability is recurrently introduced and selected against; no carrier advantage is required.	Covered inherited components show negative selection under natural fertility, with allele frequencies explained by mutation, drift, or linkage.	Favored for the maintenance arm if covered inherited components remain negative and mutation, drift, or linkage explains their frequencies.

Note. The same decision logic is applied to every row. The theory is not supported because several predictors are significant, several factors fit, or several mechanisms are biologically plausible.

A Worked Example: Same Data, Different Identification Assumptions

Two simulations are reported. They answer different questions and their numbers must not be pooled. **Simulation A** fits competing structural models to one portfolio-generated dataset in which there is no true residual mediator-outcome correlation; it asks whether explicit causal restrictions can defeat a route-free competitor. **Simulation B** sets a true residual mediator-outcome correlation and asks how badly a naive mediation analysis is degraded by it. Table 3 reports Simulation A; the confounding-degradation results are given in the text and in Figure 5A.

Table 3 *Simulation A: Competing Structural Models Fitted to the Same Portfolio-Generated Data*

Model or quantity	Causal restrictions	Value
Strict portfolio SCM	$G \rightarrow M_G \rightarrow A$ and $I \rightarrow M_I \rightarrow A$; no direct source-to-attraction paths; no residual M–A correlations	−9.870
Strict route-free SCM	G and I predict A directly; M_G and M_I have no remaining link to A	−9.922
Route-free sensitivity SCM	Direct source effects plus freely estimated residual M_G–A and M_I–A associations	−9.870
Strict-portfolio advantage over the strict route-free model	Difference in held-out log likelihood per observation	+.052
Fitted residual correlations in the sensitivity model	$\rho(e_{MG}, e_A)$; $\rho(e_{MI}, e_A)$	.225; .236

Note. SCM = structural causal model. n = 20,000; seed 20260825; 70% training and 30% hold-out; four attraction indicators with reliability .70; true standardized indirect effect .10 for each path; no true residual mediator-outcome correlation. The first three values are held-out Gaussian log likelihoods per observation; the fourth is a difference in those units; the fifth is a pair of correlations. Simulation B, in which a true residual mediator-outcome correlation was imposed, is reported in the text and in Figure 5A.

Simulation A: Model Comparison on One Dataset

With n = 20,000 and seed 20260825, independent standardized sources G and I each caused its mechanism with coefficient .40, and each mechanism caused latent attraction A with coefficient .25. The true standardized indirect effect for each path was therefore .10, exactly the materiality threshold. Four attraction indicators had reliability .70. Models were estimated in 70% of the sample and evaluated in the same 30% hold-out.

The strict portfolio structural causal model achieved a held-out Gaussian log likelihood of −9.870 per observation. A strict route-free model, in which G and I acted directly on A and the mechanism markers had no remaining link to A, fit worse at −9.922. That difference of .052 log-likelihood units per observation shows that a genuinely route-free model can be rejected when its causal restrictions are made explicit. But a route-free sensitivity model that allowed residual M_G–A and M_I–A associations fit at −9.870, essentially tying the portfolio model, with fitted residual correlations of .225 and .236.

Those two fitted values require careful interpretation, and the interpretation determines the robustness gate. They are not estimates of the residual confounding present in the data, because there was none: they arise because a model deprived of the mediator paths must absorb genuine mediation as residual covariance. What they do establish is that residual correlations of roughly .23 are observationally admissible in this design: that is the confounding a skeptic needs in order to explain away a true path of exactly the materiality threshold, and the data cannot exclude it. A stress test capped at an assumed residual correlation of .10 therefore never reaches the region in which the competing explanation lives. No gate below .23 can make a single mediation estimate immune to that skeptic, and a gate at .23 or above would require

naive effects of about .144 and so exclude true effects at the threshold itself. The gate is set at $\rho^* \geq .20$ so that a path must survive most of the admissible range, and the exposure that remains is closed by the second key, independent identification, rather than by a stricter ceiling.

Simulation B: Confounding Degradation

In a separate simulation the structural coefficients remained .40 and .25, so each causal indirect effect remained .10, but a true residual mediator-outcome correlation was imposed. At a true residual correlation of .10, a naive no-confounding mediation analysis estimated the two indirect effects at approximately .136 and .136. At .20, the corresponding estimates were approximately .167 and .168. These are not small distortions around the decision boundary.

The example changes the burden of proof. A strict route-free model can misfit the same data that generated a material path, but observational equivalence returns as soon as plausible residual mediator-outcome association is admitted. The causal-portfolio theory therefore requires the two-key rule in Table 1: confounding-robust prospective mediation plus an identification source that does not rely on sequential ignorability. Factor solutions, marginal associations, and mediation regressions alone cannot establish a route. Figure 5A reports the robustness curve for three candidate effect sizes, and Figure 5B the operating characteristics of the rejection decision.

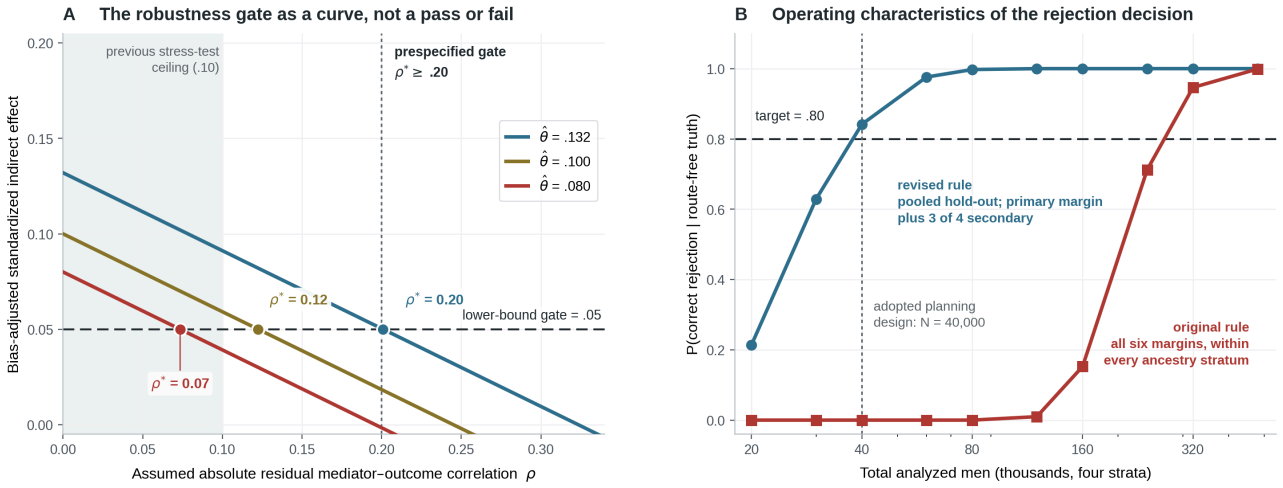


Figure 5 *The Robustness Gate as a Curve and the Operating Characteristics of the Rejection Decision*. Panel A plots the bias-adjusted standardized indirect effect against the assumed absolute residual mediator-outcome correlation ρ for three candidate effect sizes, under the linear approximation $\hat{\theta} - a\kappa\rho$ with $a = .40$ and $\kappa = 1.02$. The robustness value ρ^* is the point at which each line crosses the lower-bound gate of .05. A path with a naive indirect effect of .10 has $\rho^* = .12$; the prespecified gate of $\rho^* \geq .20$ therefore implies a naive indirect effect of about .132. The shaded band marks the assumed ceiling of .10 used in an earlier version of this framework. Panel B plots the probability of a correct rejection under a route-free data-generating process against total analyzed sample size, for the conjunctive rule, labeled original in the panel (equivalence on every margin within every stratum), and the pooled rule adopted here, labeled revised (primary causal margin plus at least three of four secondary margins). The conjunctive rule is effectively unable to fire at any feasible design.

Empirical Constraints on the Causal Graph

Inherited Sources: Polygenic Liability and Maternal-Line Signatures

Twin, family, linkage, and association studies establish familial and inherited influence without high penetrance. Concordance and heritability estimates vary with phenotype, recruitment, prevalence assumptions, nonresponse, and model specification. That variation is compatible with causal heterogeneity but is not evidence for it by itself. The invariant conclusion is narrower: inherited liability is distributed, probabilistic, and accompanied by large individual-specific variation (Bailey & Pillard, 1991; Bailey et al., 1999, 2000; Kendler et al., 2000; Långström et al., 2010).

Genome-wide results support highly distributed effects and permit an omnigenic architecture in which many variants influence attraction through broadly expressed regulatory networks (Boyle et al., 2017). Broad same-sex behavior shares genetic signal with attraction, fantasy, and identity, but the measures are not equivalent; the same GWAS also showed differences between ever-versus-never behavior and proportion of same-sex partners (Ganna et al., 2019, 2021; Hamer et al., 2021). Reported common-SNP heritability for same-sex sexual behavior ranged from approximately .08 to .25 depending on phenotype and sample; the planning value adopted for repeated attraction is .15, and the actual denominator for the coverage criterion must be estimated in the discovery sample rather than assumed. Functional or genomic partitions can summarize common-SNP heritability without pretending that each partition is a sparse causal factor (Finucane et al., 2015). The theory predicts substantial shared variance between attraction and behavioral propensity while exposing their residual separation to direct genomic test.

Xq28 is a contested source signature, not a route. Selected brother-pair studies reported excess sharing, an independent Canadian study did not support a large effect, unselected family histories did not show the expected general maternal-line excess, and later linkage scans produced ascertainment-sensitive results (Bailey et al., 1999; Hamer et al., 1993; S. Hu et al., 1995; Mustanski et al., 2005; G. Rice et al., 1999; Sanders et al., 2015). Extreme maternal X-inactivation skewing and the Han Chinese associations at FMR1NB and ZNF536 remain biologically relevant leads, but neither establishes one maternally inherited path (Bocklandt et al., 2006; S.-H. Hu et al., 2021). A valid inherited path requires a replicated molecular signal, a mechanism linking it to prenatal development, and predicted family and fitness consequences in held-out samples.

Prenatal Endocrine, Placental, and Regulatory Mechanisms

The organizational-activational principle supplies the strongest general mechanism for the sexual differentiation of partner preference. Early steroid action alters gene expression and neural organization, whereas adult hormones primarily regulate libido and activation rather than reverse the target of attraction (Hines, 2011; Phoenix et al., 1959). Swift-Gallant et al. (2023)

substantially strengthen this part of the theory. Across experiments of nature and converging biomarkers, the review concludes that early gonadal hormone action influences human sexual orientation, with particularly strong evidence that androgens increase gynephilia in people raised as girls and suggestive evidence that estrogens increase androphilia. That evidence does not by itself establish the direction or magnitude of ordinary variation among typically developing males.

The same review also prevents an overstatement. Evidence for within-male variation is mixed: men with typical external anatomy may have passed a threshold sufficient for genital masculinization while differing in tissue-specific sensitivity, developmental timing, immune exposure, or nonlinear hormone effects. Swift-Gallant et al. explicitly consider multiple male pathways and the possibility that both relatively low and relatively high androgenic conditions could be associated with different nonheterosexual subgroups. The causal-portfolio theory therefore locates the bottleneck at prenatal organization but rejects a scalar deficit model.

Placental transport and metabolism are causal mechanisms in their own right because fetal tissue exposure is not identical to maternal serum concentration. Receptor expression, local steroid conversion, co-regulator availability, fetal genotype, maternal genotype, and cell-specific regulatory state can redirect or buffer the same circulating signal. Regulatory variation is not an unexplained residual class. It is a set of measurable mechanisms that must predict tissue-relevant transcription, interact with upstream exposures, and survive prospective and functional replication.

Canalization is used in its conventional sense of developmental robustness. Thresholding, monozygotic discordance, epigenetic persistence, and stochastic divergence can occur within a robust system, but none demonstrates canalization by itself. The transgenerational epigenetic model of W. R. Rice et al. (2012) remains a candidate implementation rather than a required premise. Adult blood or saliva methylation classifiers (Ngun et al., 2015) cannot reconstruct a prenatal neural state, and reprogrammed cells undergo extensive epigenetic resetting while retaining memory of the adult source tissue rather than of any prenatal neural state (Kim et al., 2010). A regulatory claim becomes causal only when prospective prenatal measures, identified mediation, and perturbation converge.

Maternal Immune and Pregnancy-History Mechanisms

The fraternal birth-order association is the strongest evidence for a pregnancy-dependent influence not reducible to the son's inherited orientation liability. Many studies report higher odds of male homosexuality with additional older biological brothers, and the maternal immune hypothesis proposes progressive immunization to male-specific fetal proteins (Blanchard, 2004, 2018; Blanchard & Bogaert, 1996; Blanchard & Klassen, 1997; Bogaert, 2006). Bogaert et al. (2018, 2022) reported the predicted pattern for selected NLGN4Y sequences, but

the study was cross-sectional, used maternal samples decades after pregnancy, and required a sequence-specific correction. It is a mechanism-linked lead, not a demonstrated mediator.

Older-brother count is an exposure proxy rather than a route indicator. Population-attributable fractions near 15% and 29% depend on different exposure definitions and assumptions and do not assign individual cases (Blanchard & Bogaert, 2004; Cantor et al., 2002). The current logistic approach of Blanchard et al. (2025) separates older-brother, older-sister, and sibship-size contrasts under explicit ceteris-paribus substitutions and recommends theory-neutral labels for the resulting coefficients. Dutch register results found birth-order effects not confined to men, and simulations show that some family-size controls can induce or exaggerate an association (Ablaza et al., 2022; Vilsmeier et al., 2023). The immune path therefore requires serial within-woman immunization, index-pregnancy fetal exposure, prospective prediction of adult attraction, and mediation of a material part of the biological-older-brother association across a preregistered antigen and cellular panel.

Developmental Intermediates, Behavior, and Comparative Theory

The graph includes sex discrimination because comparative work shows that same-sex sexual behavior can arise from mating decisions without same-sex attraction. N. W. Bailey and Zuk (2009) established the general evolutionary distinction, and Lerch (2026) connects it to classic mate-choice theory: imperfect sex discrimination, acceptance thresholds, encounter rates, opportunity costs, mating effort, and social context can generate or maintain same-sex behavior. Lerch explicitly distinguishes same-sex attraction from behavior and does not claim that the animal framework directly explains human sexuality. Its importance here is causal: it supplies a competing path from general mating strategy to behavior that bypasses enduring male-directed valuation.

This distinction is also visible within the human genomic record. Risk-taking can mediate a positive fertility association for bisexual behavior while exclusive same-sex behavior remains negatively associated with reproduction (Song & Zhang, 2023, 2024). A behavioral-propensity component may therefore have a different evolutionary history from the component that produces high $\bar{A}_{\mathrm{M}}$ and low $\bar{A}_{\mathrm{F}}$. Critical Study 2 asks whether that difference survives reliable attraction measurement and multivariate genomic modeling; it does not assume the answer.

Comparative evidence sets boundaries rather than supplying a direct human explanation. Same-sex behavior is phylogenetically widespread, plastic, and sometimes socially functional. Comparative reconstructions support repeated gains, losses, and associations with sociality rather than one assumed ancestral state (Gómez et al., 2023; Monk et al., 2019). In male rhesus macaques it was common, repeatable, modestly heritable, and not detectably associated with reduced reproduction in one intensively studied population (Clive et al., 2023). Male-oriented rams are more directly relevant because a stable minority chooses males even when estrous females are available, demonstrating that enduring male-oriented partner preference can arise

in a mammal without human identity or culture (Roselli et al., 2011). None of these findings establishes homology with human androphilia.

Neurobiological Convergence Without a Unitary Brain Type

The expected common endpoint is functional, not gross anatomical. Interacting hypothalamic, striatal, oxytocin, vasopressin, dopamine, attachment, and social-reward systems are plausible substrates of cue valuation and pair bonding (Blumenthal & Young, 2023). Different upstream paths can converge on relative incentive value without producing one homosexual brain type.

The postmortem and imaging record is too inconsistent to define a route. INAH3 is sexually dimorphic, but the best-controlled replication did not confirm an orientation difference in volume or neuron number; the anterior commissure and suprachiasmatic findings remain unreplicated or were replaced by different molecular observations; and HIV/AIDS ascertainment confounded most homosexual postmortem samples (Allen & Gorski, 1992; Byne et al., 2001; S.-H. Hu et al., 2021; LeVay, 1991; Swaab & Hofman, 1990). A systematic review found conflicting results, extensive overlap, small samples, and no consistent phenotype (Frigerio et al., 2021). A larger UK Biobank imaging-genetics analysis detected small multivariate cross-sex shifts, predominantly in women, and regional associations with a same-sex-behavior polygenic score; its broad behavioral phenotype and modest effects do not identify a prenatal path (Abé et al., 2021). Route-stratified neurobiology becomes informative only after upstream causal measures define the strata.

Evolutionary Accounting

Inherited Components

For inherited allele or molecular component A_j with frequency $p_{j,t}$, separate selection from mutation. The selection-only change is $\Delta p_{sel} = p_j(1 - p_j)\,\bar{s}_j(p_{j,t})$. Balancing selection is admissible only when $\bar{s}_j(p_j^*) = 0$ at an internal frequency and the selection-only recurrence is locally stable. With reversible mutation, $\Delta p_{full} = \Delta p_{sel} + \mu_j(1 - p_j) - \nu_j\, p_j$. If $\bar{s}_j$ is negative at every frequency, mutation can still create an internal equilibrium, but that equilibrium is mutation-selection balance, not balancing selection. Antagonistic pleiotropy, sexual antagonism, and overdominance predict different dominance, chromosomal, frequency, and sex-specific signatures (Gavrilets & Rice, 2006).

The carrier ledger includes heterosexual men, androphilic men, women with different attraction profiles, and every other phenotype through which the component alters reproduction. Positive effects on partner number, paternal care, mate appeal, or female fertility matter only if their carrier frequencies and effect sizes offset the cost in expressed androphilic males. Miller's (2000) threshold arithmetic shows how a small advantage in numerous

nonexpressing carriers could offset a large cost in a rare high-liability configuration under its illustrative allele-frequency assumptions, but the arithmetic is not evidence that the required coupling exists.

Current evidence is mixed. Heterosexual men with nonheterosexual twins reported more opposite-sex partners, and broad same-sex-behavior signal correlates genetically with opposite-sex partner number (Zietsch et al., 2008, 2021). Risk-taking explains the advantageous association for bisexual behavior in contemporary data (Song & Zhang, 2024). The desirable-dad hypothesis predicts a pathway through paternal care and mate value but lacks direct natural-fertility confirmation (Felesina & Zietsch, 2024). Modern partner number is not historical reproductive success, and a broad behavior score is not a score for enduring androphilia.

The Explicit Covered-Portfolio Commitment

The compensatory-maintenance claim is deliberately stronger than an existential search for one favorable component. Let the full common-variant signal be divided into mutually exclusive, preregistered genome partitions: functional annotations, nonoverlapping linkage-disequilibrium blocks, or disjoint score bins. A partition may contain thousands of correlated effects and is not assumed to be an orthogonal biological factor. A partition is selection-valid only after out-of-sample heritability or score replication, within-family validation, measurement invariance, ancestry-specific linkage-disequilibrium calibration, and transport to a natural-fertility regime. Let C be the fraction of out-of-sample SNP heritability covered by selection-valid partitions. The coverage and compensation criteria are given in Table 1.

For the covered partitions, let K_C be the expected reproductive cost attributable to expression in androphilic males under natural fertility and G_C the carrier-weighted gain in all other male and female contexts. No selection effect is imputed to the uncovered fraction $1 - C$. Sensitivity analyses may display hypothetical residual effects across an explicitly labeled range, but those assumptions do not enter the verdict and cannot turn an extrapolation into a bound.

Modern polygenic scores cannot simply be carried into historical populations. Critical Study 5 requires ancestry- and cohort-specific linkage-disequilibrium calibration, phenotype measurement invariance, within-family validation, direct genotyping in the earliest available natural-fertility cohorts, explicit gene-by-environment and fertility-regime interaction, and modeling of assortative mating. Ancestral imputation is secondary sensitivity analysis. Feasibility is stated rather than hidden: no existing dataset combines a genome-wide study of repeated attraction with all of these gates and directly genotyped natural-fertility outcomes. Planning simulations must therefore examine discovery samples of 250,000, 500,000, and 1,000,000 or more participants, followed by natural-fertility validation resources, and select the lowest design that attains the coverage and precision criteria.

The adverse covered-portfolio pattern is exact. The compensatory-maintenance claim for the covered majority fails when the coverage and transport criteria are met, the one-sided upper bound on aggregate selection across the covered partitions remains below zero, and a fitted mutation-selection, drift, or linkage process reproduces their frequencies without compensation; the noncompensatory account then replaces compensation for that majority (Table 2). If the covered aggregate is negative but no noncompensatory process reproduces the observed frequencies, the covered ledger is reported as unbalanced: compensation is not demonstrated, and the discrepancy between measured selection and observed frequency is not read as evidence of unmeasured compensation. Developmental heterogeneity could survive either result. Compensation for the selection-valid covered majority would not.

Recurrent Prenatal Exposures and Developmental Constraint

A pregnancy-history-dependent immune event can recur in every generation without orientation being vertically inherited. Its incidence depends on the distribution of prior male gestations and the conditional probability that a specified immune process contributes after each history. This solves an incidence question, not the evolutionary question of why the causal immune specificity, placental transfer, fetal target, or downstream sensitivity persists.

The theory therefore requires a node-specific trade-off. A modification that blocks the causal vulnerability must reduce expected fitness through impaired immunity, placental function, ordinary neural differentiation, or another identified process by more than it gains through preventing the rare outcome. Generic statements that immunity or developmental robustness is useful are inadmissible. If targeted perturbation removes the relevant effect without a measurable cost, the constraint claim fails and that mechanism is deleted.

Maternal-Line Fecundity, Kin Investment, and Social Effects

Early pedigree studies reported higher fecundity among mothers and maternal female relatives, and multilocus sexually antagonistic models showed that such a pattern could maintain a low-frequency liability (Camperio Ciani & Pellizzari, 2012; Camperio Ciani et al., 2004, 2008; Iemmola & Camperio Ciani, 2009). The pooled evidence now weighs against a general female-fecundity effect: common-method analyses, a decade of pooled Samoan data, Dutch registers, Czech data, and unselected pedigrees have been unfavorable (Ablaza et al., 2022; Bailey et al., 1999; Blanchard et al., 2020; Fořt et al., 2024; Semenyna et al., 2023). A rare marker-defined sex-antagonistic component remains testable, but it is not an established pillar and disappears if marker-selected families show no predicted female or heterosexual-carrier benefit.

Kin investment is subsidiary. Samoan fa'afafine report elevated avuncular willingness, whereas a Western community study found no general family-directed resource transfer and greater estrangement on some measures (Bobrow & Bailey, 2001; Vasey & VanderLaan, 2010). The inclusive-fitness requirement is quantitative: actual transfers must raise recipient survival and reproduction enough to offset forgone direct reproduction. Questionnaire differences, coalition

correlations, and altered mating markets do not satisfy that criterion. They become maintenance mechanisms only when a genetic pathway to reproductive effects is demonstrated.

Figure 3 places the inherited and recurrent accounting problems side by side and states the coverage, covered-portfolio, and constraint gates.

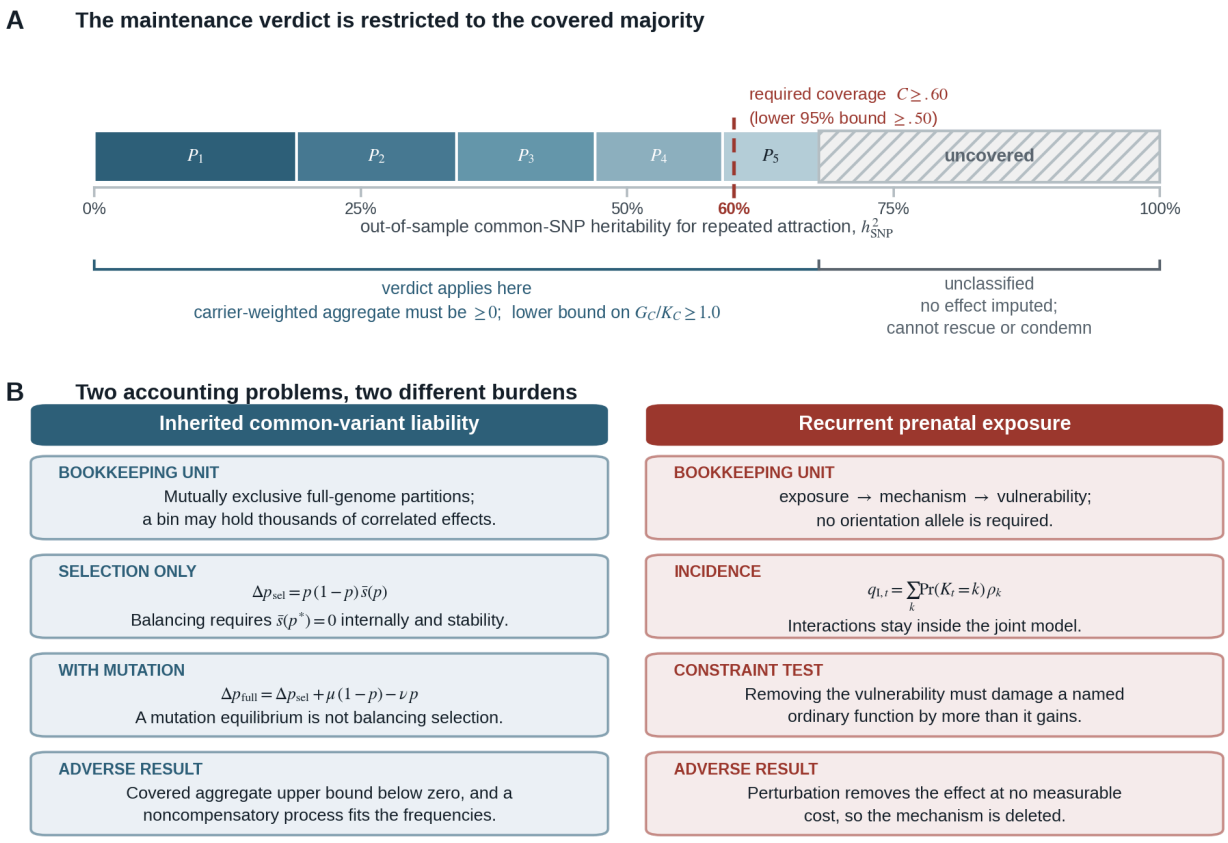


Figure 3 *Two Evolutionary Accounting Problems and the Explicit Inherited-Maintenance Commitment*. Panel A shows the coverage gate. Partitions are mutually exclusive and full-genome, and a bin may contain thousands of correlated effects; the maintenance verdict is confined to the selection-valid covered majority, and no effect is imputed to the uncovered fraction. Panel B places the two accounting problems side by side. Balancing selection is defined by the selection-only zero crossing, so a mutation equilibrium is not balancing selection. The recurrent arm requires a named exposure-to-mechanism-to-vulnerability chain and a demonstrated cost of removing it.

Critical Tests: How the Theory Can Fail

The critical studies are not a collection of desirable future projects. Each places a stated claim at risk and defines what follows from a negative result. Full sampling plans, the end-to-end simulation requirement, measurement requirements, and adverse decision rules appear in the Supplementary Material. All thresholds and margins referenced below are those in Table 1. The hierarchy is deliberately asymmetric: one result rejects the causal-portfolio core, two reject

major subtheories, others delete a candidate path or mechanism, and assay-specific nulls remain assay-specific.

Critical Study 1: Do Mechanism-Distinct Causal Paths Exist?

Claims at risk: P1, with D3 and D4 as supporting commitments. Use probability-based screening in four genetic-similarity strata and retain inclusion probabilities for every screened and enrolled participant. Oversample the region R for precision, but carry inverse-probability or calibration weights into structural, predictive, and prevalence analyses. Begin the preregistered grid at 10,000 analyzed men per stratum, the planning design derived above, and adjust in either direction as the end-to-end simulation requires. A candidate path passes only if it clears the materiality-and-identification gate in Table 1. The core is rejected only under the pooled decision rule stated above, with stratum-specific results reported as heterogeneity diagnostics.

Critical Study 2: Are Attraction and General Behavioral Propensity Genetically Distinguishable?

Claim at risk: P2. Conduct a biobank-scale multivariate genomic study with reliable independent measures of $\bar{A}_{\mathrm{M}}$, $\bar{A}_{\mathrm{F}}$, fantasy, identity, behavior, partner proportion, sociosexuality, risk-taking, openness, mating effort, and reproductive outcomes in both sexes. Compare one-factor and layered genomic structural models with measurement error and cross-sex covariance modeled explicitly. The genetic-layer claim fails if one factor is decisively preferred, attraction and behavior are predictively equivalent in held-out samples, and no valuation-specific heritable component survives the prespecified equivalence bound. That result does not make orientation, behavior, identity, and partnership the same phenotype; it removes the proposed two-layer genetic architecture.

Critical Study 3: Is Prenatal Organization the Principal Causal Stage?

Claim at risk: P3. Follow a prospective pregnancy cohort with fixed gestational measurement windows for steroid concentrations, receptor sensitivity, placental transport and metabolism, immune activity, fetal genotype, and repeated developmental outcomes through adult attraction. The primary analysis compares causal models in which prenatal measures mediate inherited and maternal effects against models in which equally reliable postnatal variables dominate after temporal and familial confounding are addressed. The consequential bound is numbered rather than named: the proposition fails if every adequately measured prenatal mechanism has a standardized effect below .10 with a lower 90% bound below .05 in every powered path stratum, while postnatal predictors produce robust, temporally ordered, out-of-sample effects that clear the same bound. An average null is insufficient when opposing nonlinear path effects were not measured or powered.

Critical Study 4: Does a Pregnancy-History-Dependent Maternal Immune Path Exist?

Claim at risk: the named gestational immune path. Enroll women from first pregnancy, collect serial samples in fixed windows of every subsequent pregnancy, assay corrected NLGN4Y sequences and a preregistered broad Y-antigen and cellular panel, measure placental transfer and fetal exposure, and follow sons to repeated $\bar{A}_{\mathrm{M}}$ and $\bar{A}_{\mathrm{F}}$ measures. The specified NLGN4Y path fails if its serial and prospective predictions fail. The broader immune path is deleted if no male-gestation-specific immunization occurs across the preregistered panel, index-pregnancy exposure does not predict attraction, and the mechanism does not mediate a material fraction of the older-brother association within precise bounds.

Critical Study 5: Was the Covered Inherited Portfolio Nonnegative Under Natural Fertility?

Claim at risk: P4. Apply the components from Critical Study 2 only after the coverage, portability, ancestry, and measurement criteria in Table 1 are met. Use directly genotyped natural-fertility cohorts, validated pedigrees, within-family analyses, both sexes, expressed and nonexpressed carriers, completed reproduction, survival of descendants, assortative mating, and fertility-regime interaction. Estimate K_{C}, G_{C}, $G_{\mathrm{C}}/K_{\mathrm{C}}$, partition-specific and aggregate covered selection, and allele-frequency predictions. The compensatory claim for the covered majority is rejected when its one-sided aggregate upper bound is below zero and a fitted noncompensatory process explains the observed frequencies; a negative aggregate that no fitted process explains is reported as an unbalanced ledger, not as compensation. Uncovered partitions remain unclassified; a nonportable score or modern partner-count association cannot alter the verdict.

Critical Study 6: Does a Named Placental or Regulatory Mechanism Carry a Causal Path?

Claim at risk: a specified placental or regulatory implementation. Combine prospective placental, cord, and neonatal measures with genetically informed designs and targeted perturbation in tissue-relevant cellular or animal systems. The named mechanism must temporally predict attraction-relevant development, mediate an upstream source effect, interact as predicted with another path when interaction is claimed, and change a relevant developmental endpoint when perturbed. It is deleted if all four requirements fail within prespecified bounds. A null adult blood classifier or a null after induced-pluripotent reprogramming rejects only that assay because neither reconstructs prenatal neural regulation.

Figure 4 separates rejection of the core from rejection of genetic layering, covered-portfolio compensation, a prenatal path, or one assay.

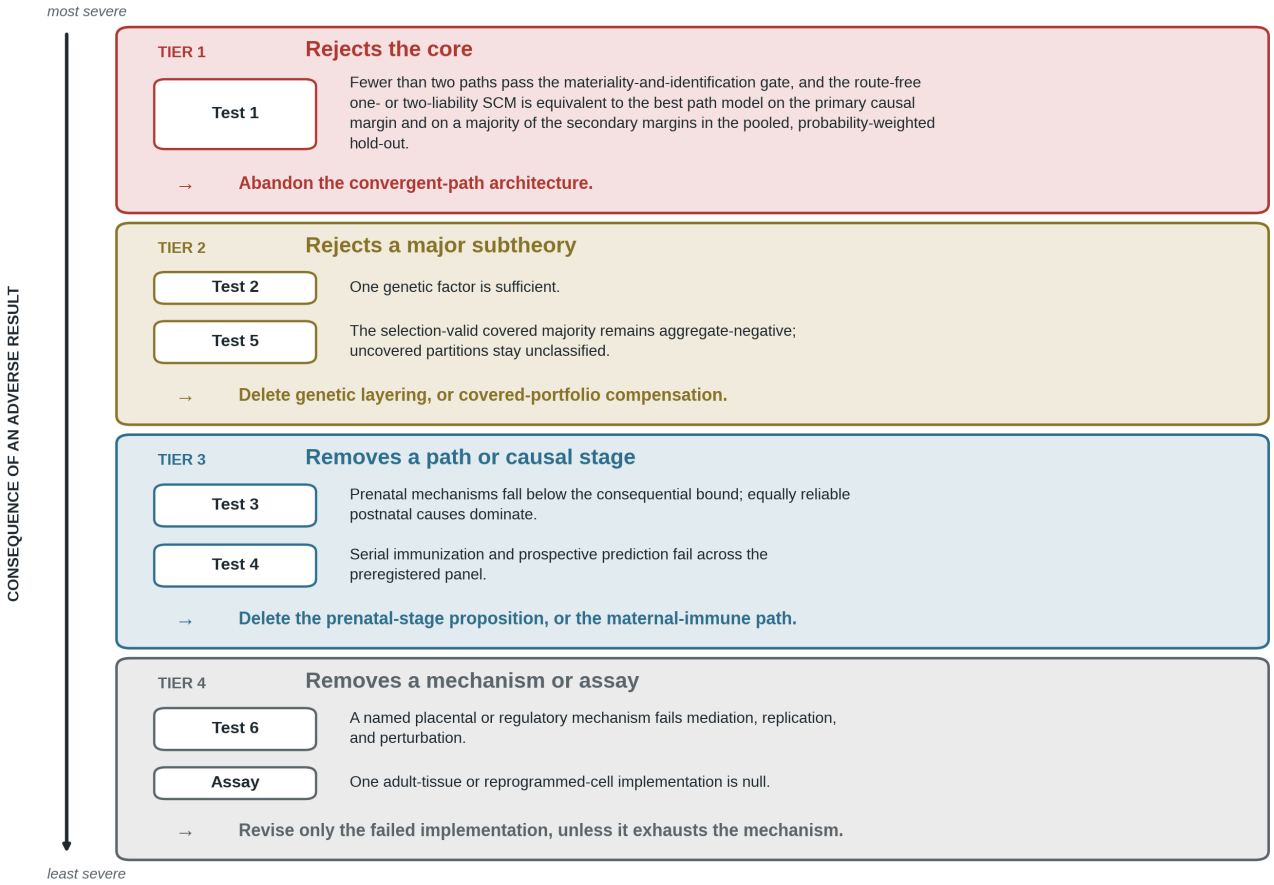


Figure 4 *The Critical Studies Arranged by the Consequence of an Adverse Result*. Tiers are ordered by the consequence of an adverse result; Tests 1 through 6 in the panel are Critical Studies 1 through 6. Critical Study 1 rejects the convergent-path core only when fewer than two paths pass the materiality-and-identification gate and the route-free comparator is equivalent on the primary causal margin and a majority of the secondary margins in the pooled, probability-weighted hold-out. Critical Study 2 rejects genetic layering; Critical Study 5 rejects compensation for the selection-valid covered majority. Critical Studies 3 and 4 remove a causal stage or path. Critical Study 6 and peripheral assay nulls remove only the named implementation unless it exhausts the mechanism.

Evolutionary Synthesis

The theory makes one decisive causal declaration: enduring male androphilia is not merely a high score on one or two route-free liabilities. It is a convergent attraction phenotype to which at least two mechanism-distinct prenatal paths make material contributions, in the precise sense given in Table 1. Those paths may share inherited sources, interact within the same pregnancy, and converge onto the same smooth attraction distribution. They are distinguished by causal restrictions and identified interventions, not by circles, clusters, factor labels, or a mediation regression alone.

The theory makes a second decisive evolutionary declaration: no single maintenance process carries the entire phenotype. Inherited common-variant liability is tested through selection-

valid full-genome partitions and allele-level accounting across both sexes. The verdict is confined to the covered majority; uncovered partitions are not assigned the effect distribution of the covered set. Recurrent prenatal exposure faces an incidence and constraint problem. Observed prevalence is generated by their joint distribution. A candidate path that lacks its predicted robust causal evidence, compensation, or constraint is deleted rather than absorbed into an unspecified remainder.

Recent work sharpens rather than dilutes the theory. Swift-Gallant et al. (2023) strengthen prenatal organization while showing why male variation cannot be reduced to a single monotonic androgen scale. Lerch (2026) shows, through classic mate-choice theory, how same-sex behavior can emerge from sex discrimination, acceptance thresholds, and mating strategy without same-sex attraction. Together these findings require the architecture to separate prenatal mechanisms from developmental intermediates and attraction from behavior.

The framework is therefore more than an umbrella only insofar as its additional predictions survive. It must outperform multidimensional continuous liability, VanderLaan-style profile heterogeneity, broad polygenicity plus measurement, independent mechanisms without integrated accounting, and mutation-selection explanations without carrier benefit. The comparison, the indeterminate regions, and the outcomes that would defeat the theory are all explicit, and mediation alone is explicitly denied the power to establish a route.

Conclusion

Male androphilia is a minority endpoint of sexually differentiated development whose evolutionary persistence must be explained at the level of its causal components. Offspring genotype, maternal genotype and physiology, pregnancy exposure, and stochastic variation act through prenatal endocrine, placental, immune, and regulatory mechanisms. These mechanisms alter sex discrimination, cue-specific incentive salience, attachment, and social reward, producing repeated male-directed and female-directed attraction. Behavior, identity, and partnership follow from that attraction under additional motivational and social constraints; they are not substitutes for it.

The conventional paradox is not dissolved by saying that causes differ. It is replaced by a stricter set of questions. Which directed paths are causal? Which effects survive explicit mediator-outcome confounding sensitivity and independent identification? Which edge-specific effects are identified, and which require joint interventions or bounds? Which selection-valid covered partitions were nonnegative under natural fertility? Which recurrent vulnerabilities are constrained or costly to remove?

The causal-portfolio theory answers with five noncollapsed levels, a numbered materiality threshold, a two-key path-identification rule, a robustness gate reported as a curve, an explicit distinction between node and path effects, a single scale on which every margin is stated, a

majority-coverage criterion without extrapolation to the uncovered fraction, and adverse results that compel abandonment or deletion. Fewer than two qualifying paths plus route-free equivalence under the pooled decision rule rejects the core. A negative selection-valid covered portfolio rejects compensatory maintenance for that majority. Failure of a named immune, endocrine, placental, or regulatory mechanism removes that path.

Two limits are stated plainly rather than absorbed. The rejection decision has adequate operating characteristics only at sample sizes near 40,000 analyzed men, and the corresponding recruitment burden runs to millions of invitations. Critical Study 5 cannot presently be run at all. The theory is falsifiable in the specific sense that its adverse patterns and their error rates are prespecified, not in the looser sense that a decisive study is currently available. Nothing here is protected by the word heterogeneity.

Declarations

Funding. No funding was received to assist with the preparation of this manuscript. No funds, grants, or other support were received for conducting this work.

Competing interests. The author has no competing interests to declare that are relevant to the content of this article. The author operates MooreMetrics, the independent website on which this manuscript is published; it is otherwise unrelated to the subject matter of the article.

Ethics approval. Not applicable. This is a theoretical article. It reports no new data collected from human participants or animals. All empirical findings discussed are drawn from the published literature, and all quantitative results reported by the author are derived from simulated data.

Availability of data and material. No new empirical data were generated. The simulated datasets underlying the worked example, the robustness curve, and the decision-rule operating characteristics are fully specified by the generating parameters and random seeds reported in the article and can be reproduced from them.

Code availability. The simulation and figure-generation scripts are available from the author on request.

Use of generative AI and AI-assisted technologies. The author used Claude Opus 4.8, GPT-5.6 Sol Pro, and Grok 4.3 to prepare initial drafts and for proofreading. The author reviewed and edited all content and takes full responsibility for it.

References

Abé, C., Lebedev, A., Zhang, R., Jonsson, L., Bergen, S. E., Ingvar, M., Landén, M., & Rahman, Q. (2021). Cross-sex shifts in two brain imaging phenotypes and their relation to polygenic scores for same-sex sexual

behavior: A study of 18,645 individuals from the UK Biobank. *Human Brain Mapping*, *42*, 2292–2304. https://doi.org/10.1002/hbm.25370

Ablaza, C., Kabátek, J., & Perales, F. (2022). Are sibship characteristics predictive of same sex marriage? An examination of fraternal birth order and female fecundity effects in population-level administrative data from the Netherlands. *The Journal of Sex Research*, *59*, 671–683. https://doi.org/10.1080/00224499.2021.1974330

Allen, L. S., & Gorski, R. A. (1992). Sexual orientation and the size of the anterior commissure in the human brain. *Proceedings of the National Academy of Sciences of the United States of America*, *89*, 7199–7202. https://doi.org/10.1073/pnas.89.15.7199

Apostolou, M. (2022). The direct reproductive cost of same-sex attraction: Evidence from two nationally representative U.S. samples. *Archives of Sexual Behavior*, *51*, 1857–1864.

Avin, C., Shpitser, I., & Pearl, J. (2005). Identifiability of path-specific effects. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence* (pp. 357–363).

Bailey, J. M., & Pillard, R. C. (1991). A genetic study of male sexual orientation. *Archives of General Psychiatry*, *48*, 1089–1096.

Bailey, J. M., Dunne, M. P., & Martin, N. G. (2000). Genetic and environmental influences on sexual orientation and its correlates in an Australian twin sample. *Journal of Personality and Social Psychology*, *78*, 524–536. https://doi.org/10.1037/0022-3514.78.3.524

Bailey, J. M., Pillard, R. C., Dawood, K., Miller, M. B., Farrer, L. A., Trivedi, S., & Murphy, R. L. (1999). A family history study of male sexual orientation using three independent samples. *Behavior Genetics*, *29*, 79–86.

Bailey, N. W., & Zuk, M. (2009). Same-sex sexual behavior and evolution. *Trends in Ecology & Evolution*, *24*, 439–446. https://doi.org/10.1016/j.tree.2009.03.014

Blanchard, R. (2004). Quantitative and theoretical analyses of the relation between older brothers and homosexuality in men. *Journal of Theoretical Biology*, *230*, 173–187.

Blanchard, R. (2018). Fraternal birth order, family size, and male homosexuality: Meta-analysis of studies spanning 25 years. *Archives of Sexual Behavior*, *47*, 1–15.

Blanchard, R., & Bogaert, A. F. (1996). Homosexuality in men and number of older brothers. *American Journal of Psychiatry*, *153*, 27–31.

Blanchard, R., & Bogaert, A. F. (2004). Proportion of homosexual men who owe their sexual orientation to fraternal birth order: An estimate based on two national probability samples. *American Journal of Human Biology*, *16*, 151–157. https://doi.org/10.1002/ajhb.20006

Blanchard, R., & Klassen, P. (1997). H-Y antigen and homosexuality in men. *Journal of Theoretical Biology*, *185*, 373–378.

Blanchard, R., Kabátek, J., & Zdaniuk, B. (2025). A current approach to logistic regression analysis of birth order and sexual orientation. *Archives of Sexual Behavior*, *54*, 4013–4022. https://doi.org/10.1007/s10508-025-03275-3

Blanchard, R., Krupp, J., VanderLaan, D. P., Vasey, P. L., & Zucker, K. J. (2020). A method yielding comparable estimates of the fraternal birth order and female fecundity effects in male homosexuality. *Proceedings of the Royal Society B: Biological Sciences*, *287*, Article 20192907. https://doi.org/10.1098/rspb.2019.2907

Blumenthal, S. A., & Young, L. J. (2023). The neurobiology of love and pair bonding from human and animal perspectives. *Biology*, *12*, Article 844.

Bobrow, D., & Bailey, J. M. (2001). Is male homosexuality maintained via kin selection? *Evolution and Human Behavior*, *22*, 361–368. https://doi.org/10.1016/S1090-5138(01)00074-5

Bocklandt, S., Horvath, S., Vilain, E., & Hamer, D. H. (2006). Extreme skewing of X chromosome inactivation in mothers of homosexual men. *Human Genetics*, *118*, 691–694. https://doi.org/10.1007/s00439-005-0119-4

Bogaert, A. F. (2006). Biological versus nonbiological older brothers and men's sexual orientation. *Proceedings of the National Academy of Sciences of the United States of America*, *103*, 10771–10774.

Bogaert, A. F., Skorska, M. N., Wang, C., Gabrie, J., MacNeil, A. J., Hoffarth, M. R., VanderLaan, D. P., Zucker, K. J., & Blanchard, R. (2018). Male homosexuality and maternal immune responsivity to the Y-linked protein NLGN4Y. *Proceedings of the National Academy of Sciences of the United States of America*, *115*, 302–306. https://doi.org/10.1073/pnas.1705895114

Bogaert, A. F., Skorska, M. N., Wang, C., Gabrie, J., MacNeil, A. J., Hoffarth, M. R., VanderLaan, D. P., Zucker, K. J., & Blanchard, R. (2022). Correction to supporting information for Bogaert et al., Male homosexuality and maternal immune responsivity to the Y-linked protein NLGN4Y. *Proceedings of the National Academy of Sciences of the United States of America*, *119*, Article e2212767119. https://doi.org/10.1073/pnas.2212767119

Boyle, E. A., Li, Y. I., & Pritchard, J. K. (2017). An expanded view of complex traits: From polygenic to omnigenic. *Cell*, *169*, 1177–1186. https://doi.org/10.1016/j.cell.2017.05.038

Byne, W., Tobet, S., Mattiace, L. A., Lasco, M. S., Kemether, E., Edgar, M. A., Morgello, S., Buchsbaum, M. S., & Jones, L. B. (2001). The interstitial nuclei of the human anterior hypothalamus: An investigation of variation with sex, sexual orientation, and HIV status. *Hormones and Behavior*, *40*, 86–92. https://doi.org/10.1006/hbeh.2001.1680

Camperio Ciani, A., & Pellizzari, E. (2012). Fecundity of paternal and maternal non-parental female relatives of homosexual and heterosexual men. *PLoS ONE*, *7*, Article e51088. https://doi.org/10.1371/journal.pone.0051088

Camperio Ciani, A., Cermelli, P., & Zanzotto, G. (2008). Sexually antagonistic selection in human male homosexuality. *PLoS ONE*, *3*, Article e2282. https://doi.org/10.1371/journal.pone.0002282

Camperio Ciani, A., Corna, F., & Capiluppi, C. (2004). Evidence for maternally inherited factors favouring male homosexuality and promoting female fecundity. *Proceedings of the Royal Society B: Biological Sciences*, *271*, 2217–2221.

Cantor, J. M., Blanchard, R., Paterson, A. D., & Bogaert, A. F. (2002). How many gay men owe their sexual orientation to fraternal birth order? *Archives of Sexual Behavior*, *31*, 63–71. https://doi.org/10.1023/A:1014031201935

Cinelli, C., & Hazlett, C. (2020). Making sense of sensitivity: Extending omitted variable bias. *Journal of the Royal Statistical Society: Series B*, *82*, 39–67. https://doi.org/10.1111/rssb.12348

Clive, J., Flintham, E., & Savolainen, V. (2023). Same-sex sociosexual behaviour is widespread and heritable in male rhesus macaques. *Nature Ecology & Evolution*, *7*, 1287–1301. https://doi.org/10.1038/s41559-023-02111-y

Felesina, T. (2026). Understanding the evolutionary history of human male and female same-sex sexual behavior requires a greater appreciation of its shared genetic architecture. *Archives of Sexual Behavior*, *55*, 1451–1453. https://doi.org/10.1007/s10508-026-03478-2

Felesina, T., & Zietsch, B. P. (2024). The desirable dad hypothesis: Male same-sex attraction as the product of selection for paternal care via antagonistic pleiotropy. *Archives of Sexual Behavior*, *53*, 1731–1745.

Finucane, H. K., Bulik-Sullivan, B., Gusev, A., Trynka, G., Reshef, Y., Loh, P.-R., Anttila, V., Xu, H., Zang, C., Farh, K., Ripke, S., Day, F. R., Purcell, S., Stahl, E., Lindstrom, S., Perry, J. R. B., Okada, Y., Raychaudhuri, S., Daly, M. J., … Price, A. L. (2015). Partitioning heritability by functional annotation using genome-wide association summary statistics. *Nature Genetics*, *47*, 1228–1235. https://doi.org/10.1038/ng.3404

Fořt, J., Flegr, J., Kuba, R., & Kaňková, Š. (2024). Fertility of Czech gay and straight men, women, and their relatives: Testing the sexually antagonistic gene hypothesis. *Archives of Sexual Behavior*, *53*, 1747–1761.

Frigerio, A., Ballerini, L., & Valdés Hernández, M. (2021). Structural, functional, and metabolic brain differences as a function of gender identity or sexual orientation: A systematic review of the human neuroimaging literature. *Archives of Sexual Behavior*, *50*, 3329–3352. https://doi.org/10.1007/s10508-021-02005-9

Ganna, A., Verweij, K. J. H., Nivard, M. G., Maier, R., Wedow, R., Busch, A. S., Abdellaoui, A., Guo, S., Sathirapongsasuti, J. F., 23andMe Research Team, Lichtenstein, P., Lundström, S., Långström, N., Auton, A., Harris, K. M., Beecham, G. W., Martin, E. R., Sanders, A. R., Perry, J. R. B., … Zietsch, B. P. (2019). Large-scale GWAS reveals insights into the genetic architecture of same-sex sexual behavior. *Science*, *365*, Article eaat7693. https://doi.org/10.1126/science.aat7693

Ganna, A., Verweij, K. J. H., Nivard, M. G., Maier, R., Wedow, R., Busch, A. S., Abdellaoui, A., Guo, S., Sathirapongsasuti, J. F., 23andMe Research Team, Lichtenstein, P., Lundström, S., Långström, N., Auton, A., Harris, K. M., Beecham, G. W., Martin, E. R., Sanders, A. R., Perry, J. R. B., … Zietsch, B. P. (2021). Response to comment on Large-scale GWAS reveals insights into the genetic architecture of same-sex sexual behavior. *Science*, *371*, Article eaba5693. https://doi.org/10.1126/science.aba5693

Gavrilets, S., & Rice, W. R. (2006). Genetic models of homosexuality: Generating testable predictions. *Proceedings of the Royal Society B: Biological Sciences*, *273*, 3031–3038. https://doi.org/10.1098/rspb.2006.3684

Gómez, J. M., González-Megías, A., & Verdú, M. (2023). The evolution of same-sex sexual behaviour in mammals. *Nature Communications*, *14*, Article 5719. https://doi.org/10.1038/s41467-023-41290-x

Hamer, D. H., Hu, S., Magnuson, V. L., Hu, N., & Pattatucci, A. M. L. (1993). A linkage between DNA markers on the X chromosome and male sexual orientation. *Science*, *261*, 321–327.

Hamer, D. H., Mustanski, B., Sell, R., Sanders, S. A., & Garcia, J. R. (2021). Comment on Large-scale GWAS reveals insights into the genetic architecture of same-sex sexual behavior. *Science*, *371*, Article eaba2941. https://doi.org/10.1126/science.aba2941

Hernán, M. A., & Robins, J. M. (2020). *Causal inference: What if*. Chapman & Hall/CRC.

Hines, M. (2011). Prenatal endocrine influences on sexual orientation and on sexually differentiated childhood behavior. *Frontiers in Neuroendocrinology*, *32*, 170–182. https://doi.org/10.1016/j.yfrne.2011.02.006

Hu, S., Pattatucci, A. M. L., Patterson, C., Li, L., Fulker, D. W., Cherny, S. S., Kruglyak, L., & Hamer, D. H. (1995). Linkage between sexual orientation and chromosome Xq28 in males but not in females. *Nature*

Genetics, *11*, 248–256. https://doi.org/10.1038/ng1195-248

Hu, S.-H., Li, H.-M., Yu, H., Liu, Y., Liu, C.-X., Zuo, X.-B., Lu, J., Jiang, J.-J., Xi, C.-X., Huang, B.-C., Xu, H.-J., Hu, J.-B., Lai, J.-B., Huang, M.-L., Liu, J.-N., Xu, D.-G., Guo, X.-C., Wu, W., Wu, X., … Xu, Y. (2021). Discovery of new genetic loci for male sexual orientation in Han population. *Cell Discovery*, *7*, Article 103. https://doi.org/10.1038/s41421-021-00341-7

Iemmola, F., & Camperio Ciani, A. (2009). New evidence of genetic factors influencing sexual orientation in men: Female fecundity increase in the maternal line. *Archives of Sexual Behavior*, *38*, 393–399. https://doi.org/10.1007/s10508-008-9381-6

Imai, K., Keele, L., & Tingley, D. (2010). A general approach to causal mediation analysis. *Psychological Methods*, *15*, 309–334. https://doi.org/10.1037/a0020761

Kendler, K. S., Thornton, L. M., Gilman, S. E., & Kessler, R. C. (2000). Sexual orientation in a U.S. national sample of twin and nontwin sibling pairs. *American Journal of Psychiatry*, *157*, 1843–1846. https://doi.org/10.1176/appi.ajp.157.11.1843

Kim, K., Doi, A., Wen, B., Ng, K., Zhao, R., Cahan, P., Kim, J., Aryee, M. J., Ji, H., Ehrlich, L. I. R., Yabuuchi, A., Takeuchi, A., Cunniff, K. C., Hongguang, H., McKinney-Freeman, S., Naveiras, O., Yoon, T. J., Irizarry, R. A., Jung, N., … Daley, G. Q. (2010). Epigenetic memory in induced pluripotent stem cells. *Nature*, *467*, 285–290. https://doi.org/10.1038/nature09342

Lakens, D., Scheel, A. M., & Isager, P. M. (2018). Equivalence testing for psychological research: A tutorial. *Advances in Methods and Practices in Psychological Science*, *1*, 259–269. https://doi.org/10.1177/2515245918770963

Långström, N., Rahman, Q., Carlström, E., & Lichtenstein, P. (2010). Genetic and environmental effects on same-sex sexual behavior: A population study of twins in Sweden. *Archives of Sexual Behavior*, *39*, 75–80. https://doi.org/10.1007/s10508-008-9386-1

Lerch, B. A. (2026). The evolution of same-sex sexual behavior: Using old theory to answer new questions. *Integrative and Comparative Biology*, *66*, Article icag015. https://doi.org/10.1093/icb/icag015

LeVay, S. (1991). A difference in hypothalamic structure between heterosexual and homosexual men. *Science*, *253*, 1034–1037.

Miller, E. M. (2000). Homosexuality, birth order, and evolution: Toward an equilibrium reproductive economics of homosexuality. *Archives of Sexual Behavior*, *29*, 1–34.

Monk, J. D., Giglio, E., Kamath, A., Lambert, M. R., & McDonough, C. E. (2019). An alternative hypothesis for the evolution of same-sex sexual behaviour in animals. *Nature Ecology & Evolution*, *3*, 1622–1631. https://doi.org/10.1038/s41559-019-1019-7

Mustanski, B. S., DuPree, M. G., Nievergelt, C. M., Bocklandt, S., Schork, N. J., & Hamer, D. H. (2005). A genomewide scan of male sexual orientation. *Human Genetics*, *116*, 272–278. https://doi.org/10.1007/s00439-004-1241-4

National Academies of Sciences, Engineering, and Medicine. (2023). *Using population descriptors in genetics and genomics research: A new framework for an evolving field*. National Academies Press. https://doi.org/10.17226/26902

Ngun, T. C., Guo, W., Ghahramani, N. M., Purkayastha, K., Conn, D., Sanchez, F. J., Bocklandt, S., Zhang, M., Ramirez, C. M., Pellegrini, M., & Vilain, E. (2015, October 8). *A novel predictive model of sexual orientation using epigenetic markers* (Abstract No. 95) [Conference presentation]. American Society of Human Genetics 2015 Annual Meeting, Baltimore, MD, United States.

Pearl, J. (2009). *Causality: Models, reasoning, and inference* (2nd ed.). Cambridge University Press.

Phoenix, C. H., Goy, R. W., Gerall, A. A., & Young, W. C. (1959). Organizing action of prenatally administered testosterone propionate on the tissues mediating mating behavior in the female guinea pig. *Endocrinology*, *65*, 369–382. https://doi.org/10.1210/endo-65-3-369

Rahman, Q., & Wilson, G. D. (2003). Born gay? The psychobiology of human sexual orientation. *Personality and Individual Differences*, *34*, 1337–1382. https://doi.org/10.1016/S0191-8869(02)00140-X

Rice, G., Anderson, C., Risch, N., & Ebers, G. (1999). Male homosexuality: Absence of linkage to microsatellite markers at Xq28. *Science*, *284*, 665–667. https://doi.org/10.1126/science.284.5414.665

Rice, W. R., Friberg, U., & Gavrilets, S. (2012). Homosexuality as a consequence of epigenetically canalized sexual development. *The Quarterly Review of Biology*, *87*, 343–368.

Robins, J. M., & Greenland, S. (1992). Identifiability and exchangeability for direct and indirect effects. *Epidemiology*, *3*, 143–155. https://doi.org/10.1097/00001648-199203000-00013

Roselli, C. E., Reddy, R. C., & Kaufman, K. R. (2011). The development of male-oriented behavior in rams. *Frontiers in Neuroendocrinology*, *32*, 164–169. https://doi.org/10.1016/j.yfrne.2010.12.007

Sanders, A. R., Martin, E. R., Beecham, G. W., Guo, S., Dawood, K., Rieger, G., Badner, J. A., Gershon, E. S., Krishnappa, R. S., Kolundzija, A. B., Duan, J., Gejman, P. V., & Bailey, J. M. (2015). Genome-wide scan demonstrates significant linkage for male sexual orientation. *Psychological Medicine*, *45*, 1379–1388. https://doi.org/10.1017/S0033291714002451

Semenyna, S. W., Gómez Jiménez, F. R., VanderLaan, D. P., & Vasey, P. L. (2023). Male androphilia, fraternal birth order, and female fecundity in Samoa: A 10-y retrospective. *Proceedings of the National Academy of Sciences of the United States of America*, *120*, Article e2313284120. https://doi.org/10.1073/pnas.2313284120

Shapley, L. S. (1953). A value for n-person games. In H. W. Kuhn & A. W. Tucker (Eds.), *Contributions to the theory of games* (Vol. 2, pp. 307–317). Princeton University Press.

Song, S., & Zhang, J. (2023). Contraception ends the genetic maintenance of human same-sex sexual behavior. *Proceedings of the National Academy of Sciences of the United States of America*, *120*, Article e2303418120. https://doi.org/10.1073/pnas.2303418120

Song, S., & Zhang, J. (2024). Genetic variants underlying human bisexual behavior are reproductively advantageous. *Science Advances*, *10*, Article eadj6958. https://doi.org/10.1126/sciadv.adj6958

Swaab, D. F., & Hofman, M. A. (1990). An enlarged suprachiasmatic nucleus in homosexual men. *Brain Research*, *537*, 141–148. https://doi.org/10.1016/0006-8993(90)90350-K

Swift-Gallant, A., Aung, T., Rosenfield, K., Dawood, K., & Puts, D. (2023). Organizational effects of gonadal hormones on human sexual orientation. *Adaptive Human Behavior and Physiology*, *9*, 344–370. https://doi.org/10.1007/s40750-023-00226-x

VanderLaan, D. P., Skorska, M. N., Peragine, D. E., & Coome, L. A. (2022). Carving the biodevelopment of same-sex sexual orientation at its joints. In D. P. VanderLaan & W. I. Wong (Eds.), *Gender and sexuality*

development: Contemporary theory and research (pp. 491–538). Springer. https://doi.org/10.1007/978-3-030-84273-4_17

VanderWeele, T. J. (2015). *Explanation in causal inference: Methods for mediation and interaction*. Oxford University Press.

VanderWeele, T. J., & Ding, P. (2017). Sensitivity analysis in observational research: Introducing the E-value. *Annals of Internal Medicine*, *167*, 268–274. https://doi.org/10.7326/M16-2607

VanderWeele, T. J., Vansteelandt, S., & Robins, J. M. (2014). Effect decomposition in the presence of an exposure-induced mediator-outcome confounder. *Epidemiology*, *25*, 300–306. https://doi.org/10.1097/EDE.0000000000000034

Vasey, P. L., & VanderLaan, D. P. (2010). Avuncular tendencies and the evolution of male androphilia in Samoan fa'afafine. *Archives of Sexual Behavior*, *39*, 821–830. https://doi.org/10.1007/s10508-008-9404-3

Vilsmeier, J. K., Kossmeier, M., Voracek, M., & Tran, U. S. (2023). The fraternal birth-order effect as a statistical artefact: Convergent evidence from probability calculus, simulated data, and multiverse meta-analysis. *PeerJ*, *11*, Article e15623.

Zietsch, B. P., Morley, K. I., Shekar, S. N., Verweij, K. J. H., Keller, M. C., Macgregor, S., Wright, M. J., Bailey, J. M., & Martin, N. G. (2008). Genetic factors predisposing to homosexuality may increase mating success in heterosexuals. *Evolution and Human Behavior*, *29*, 424–433.

Zietsch, B. P., Sidari, M. J., Abdellaoui, A., Maier, R., Långström, N., Guo, S., Beecham, G. W., Martin, E. R., Sanders, A. R., & Verweij, K. J. H. (2021). Genomic evidence consistent with antagonistic pleiotropy may help explain the evolutionary maintenance of same-sex sexual behaviour in humans. *Nature Human Behaviour*, *5*, 1251–1258.

Supplementary Material

This document supports the manuscript named above and is not intended to be read independently of it. Sections S1 through S10 give the detailed empirical record. Sections S11 and S12 define the five-level causal-path architecture, model competitors, and comparative constraints. Section S13 distinguishes node from edge interventions, supplies a worked structural-causal example fitted to one common dataset, separates selection-only balancing from mutation-selection equilibrium, and gives full protocols for six critical studies. It also states the end-to-end simulation requirement by which the sample size for Critical Study 1 is selected. Throughout, androphilia names the repeated attraction construct. When a cited study measured homosexuality, same-sex behavior, identity, partnership, or legal union, its own phenotype term is retained and is not treated as interchangeable with androphilia. Works cited here are listed in the supplementary reference list; works cited in both documents appear in both lists. Every threshold, margin, and decision rule referred to below is defined in Table 1 of the main text, which is authoritative wherever this document repeats a value.

S1 Shared Genetic Architecture and Sex-Differential Expression

The manuscript is sex-specific in its developmental explanandum, not in its population-genetic bookkeeping. Broad reviews of sexual-orientation science already require distinctions among attraction, behavior, identity, and sex-specific expression (Bailey et al., 2016). Male and female sexuality share much of the same genome, and the largest behavioral GWAS estimated a substantial cross-sex genetic correlation of approximately .63 (Felesina, 2026; Ganna et al., 2019). The relevant claim is therefore that shared inherited variation can be expressed differently across sexes, not that male and female orientation have independent genetic architectures.

The Finnish population twin study of Alanko et al. (2010) supplies the sharpest quantitative example of sex-differential expression, but its coefficients change sign between sections of the source, so the scoring directions must be stated before any of them is interpreted. In the descriptive analyses the childhood scale ran in the gender-*typical* direction, with higher scores indicating more conventional gender role behavior, and the orientation scale ran in the same-sex direction. On that scoring, gender-typical behavior correlated with adult same-sex orientation at r = −.42 in men and r = −.19 in women. Reversing the childhood scale to the atypicality direction gives +.42 and +.19: in both sexes, more gender-atypical recalled behavior is associated with more same-sex orientation, more strongly in men. These coefficients account for approximately 17.6% and 3.6% of variance, respectively. The male correlation was about 2.2 times the female correlation in magnitude, whereas the ratio of explained variances was about 4.9. The sex-by-behavior interaction was significant. Model comparisons favored additive-genetic plus nonshared-environment formulations in men and dominance or additive-plus-dominance formulations in women.

The biometric analyses used a third scoring direction, and this is where the published signs are most easily misread. For the genetic models the authors reversed the childhood scale for men so that, in both sexes, a higher score indicated more *girlish* behavior. On that girlishness scale the fitted bivariate genetic correlation with adult same-sex orientation was +.73 in men and approximately −.28 to −.30 in women. The sex difference in sign is therefore an arithmetic consequence of applying one girlishness scale to two sexes, because girlishness is gender-atypical in men and gender-typical in women. Rescored so that higher values mean gender-atypical for the participant's own sex, both correlations are positive, roughly +.73 in men and +.28 to +.30 in women. The defensible reading is a sex difference in the *magnitude* of shared genetic influence, not a reversal of its direction, and this supplement does not claim otherwise. Two further cautions belong with the estimate: the female coefficients come from different variance components in different parameterizations, an additive correlation of −.28 under an ACE model and a dominance correlation of −.30 under the better-fitting DE model, and the best-fitting model differed by sex, AE in men and DE in women, which the authors themselves note makes a direct cross-sex comparison difficult.

Evidence that partly different genetic factors operate in the two sexes comes from elsewhere in the same study, not from the sign of the bivariate correlation. The opposite-sex dizygotic twin correlation for orientation was approximately zero, well below the same-sex dizygotic values, and fixing the cross-sex additive genetic correlation to .5 produced a near-significant deterioration in fit. That is suggestive and underpowered, and it is treated here as such.

Those estimates come from one sample and cannot, by themselves, establish two architectures. The classical twin design had limited power to separate additive from nonadditive effects; orientation was measured for the past year; gender-atypical behavior was retrospective; and response was incomplete. The defensible inference is narrower: shared inherited variation is expressed with a stronger and more coherent link between early gendered behavior and adult orientation in men than in women in this dataset. A theory centered on male androphilia gains precision by modeling that sex difference while retaining female same-sex behavior and orientation in the allele-level fitness portfolio.

S2 Twin, Family, and Registry Estimates in Detail

Estimates of familial and genetic influence differ across designs. The variation is compatible with etiological mixture, but it is not uniquely diagnostic of mixture. Phenotype definition, concordance-dependent volunteer recruitment, prevalence assumptions, nonresponse, cohort and cultural differences, measurement error, model misspecification, low power, and ordinary sampling fluctuation can all change the estimates. This section reports the record without converting every difference into evidence for a biological subtype.

Selected Twin and Adoption Evidence

Bailey and Pillard (1991) reported concordance for homosexuality or bisexuality of 52% in monozygotic cotwins, 22% in dizygotic cotwins, and 11% in adoptive brothers among relatives whose orientation could be rated. Multifactorial threshold models yielded liability-heritability estimates from .31 to .74 under different assumptions about prevalence and ascertainment. Shared-environment estimates were smaller and nonsignificant, but the design did not prove the literal absence of shared environmental influence. The bimodal distribution of Kinsey scores among monozygotic cotwins was compatible with a threshold process and did not establish a major gene.

Only 9.2% of nontwin biological brothers were rated homosexual or bisexual, substantially below the dizygotic-twin concordance. Recruitment through gay publications could preferentially enroll concordant, highly visible, or strongly identified families. Such ascertainment can enrich familial loading; it does not identify which biological route produced that loading. The correct conclusion is that selected samples are useful for discovering concentrated familial signals but cannot estimate their prevalence without explicit sampling correction.

Population, Registry, and Family-History Evidence

Bailey et al. (2000) used an Australian twin-register sample rather than a homosexual-proband recruitment design. Male probandwise concordance was 20.0% for monozygotic twins and 0% for dizygotic twins under a strict Kinsey criterion, and 37.5% versus 6.3% under a broader criterion. The best-fitting univariate model assigned .45 of variance to additive genetic influence and .55 to individual-specific influence, with a point estimate of zero for shared environment. The lower concordances relative to advertisement studies demonstrate ascertainment sensitivity; they do not distinguish route mixture from recruitment and measurement effects.

In a U.S. national sample, monozygotic twins were more concordant for nonheterosexual orientation than dizygotic twins or nontwin siblings. Probandwise concordance was 31.6% for monozygotic pairs and 13.3% for same-sex dizygotic pairs, with biometric heritability estimates from .28 to .65 depending on the comparison group (Kendler et al., 2000). Confidence intervals were wide because nonheterosexual orientation was uncommon.

The Swedish Twin Registry used same-sex sexual behavior rather than enduring orientation. Among men, additive genetic influence was estimated at .39 for ever having had a same-sex partner and .34 for number of same-sex partners; individual-specific variance accounted for .61 and .66, while the point estimate for shared environment was zero (Långström et al., 2010). These are estimates for behavioral liability, not a direct decomposition of predominant or exclusive androphilia.

Bailey et al. (1999) recruited three samples designed to reduce volunteer bias: 350 nonheterosexual men consecutively admitted to an HIV clinic, 167 gay men approached at a Gay Pride event with recruitment blind to family history, and 65 men selected because they had gay brothers. In the clinic sample, 7.3% of 453 adult brothers were classified with high certainty as homosexual or bisexual, 95% CI [5%, 10%], against 3.8% of 394 sisters. In the Pride sample, 9.7% of 259 brothers were reported homosexual, 95% CI [6%, 14%], against 2.8% of 214 sisters. Against an assumed population estimate of 2.4%, the brother rates corresponded to sibling relative risks of approximately 3.0 and 4.0.

Sibling relative risks of three to four establish familial concentration under the assumptions used, but they do not logically exclude every social or environmental account. Familial prenatal exposure, correlated gene-environment processes, reporting effects, and some family-level social mechanisms can also produce recurrence. The invariant conclusion is that male same-sex orientation and behavior show familial and genetic influence without high penetrance. Recovering route structure requires direct mechanism-linked measures and explicit comparison with continuous, measurement, and ascertainment explanations; heterogeneity cannot be inferred from changing coefficients alone.

S3 The Xq28 Linkage Record and Its Ascertainment Dependence

Xq28 is treated as a contested marker of a possible maternally concentrated inherited path, not as an established subtype, a causal level, or a universal cause. The complete record includes positive selected-family studies, a direct independent null study, pedigree evidence against general maternal transmission, and later linkage scans. Figure S1 summarizes the record and the two quantitative signatures an X-influenced inherited path would require.

Hamer et al. (1993) first combined pedigree analysis in randomly ascertained homosexual probands with X-chromosome linkage in brother-pair families enriched for maternal transmission. Homosexuality was elevated among brothers, maternal uncles, and sons of maternal aunts but not among corresponding paternal relatives. In 40 homosexual brother pairs, 33 shared the distal Xq28 haplotype and seven did not, producing an estimated sharing proportion of .82. The seven discordant pairs were not a minor nuisance: they demonstrated substantial heterogeneity even within the deliberately enriched sample and prevented any inference that Xq28 was necessary or sufficient.

S. Hu et al. (1995) supplied a second selected-family test. Gay brother pairs shared Xq28 haplotypes approximately 67%–68% of the time, above the 50% expectation, while heterosexual brothers in Xq28-concordant families showed deficient sharing. No comparable linkage appeared in female sibling pairs. The authors again warned against extrapolating the result to unselected families. Together, the two early studies defined a narrow claim: a male-relevant, modest-penetrance factor near Xq28 contributes in some maternally loaded pedigrees.

Rice, Anderson, Risch, and Ebers (1999) directly tested Xq28 in 52 Canadian gay brother pairs and reported no support for a large X-linked effect at the markers examined. That study is the central independent null result and must be stated directly rather than represented only through the later technical exchange. Hamer (1999) argued that combined data still showed excess sharing and that power, ascertainment, marker coverage, and phenotyping could explain the null. Rice and colleagues replied that replication-only datasets did not significantly exceed chance under their criteria. The resulting boundary is exact: the literature neither establishes a universal Xq28 effect nor excludes a modest, sample-sensitive effect in specially ascertained families.

The strongest pedigree evidence against a general maternal-transmission model came from the same family-history study that produced the recurrence-risk estimates. Bailey et al. (1999) examined maternal versus paternal relatives in all three of their samples and found nothing resembling the predicted asymmetry. In the clinic sample, 4.4% of maternal uncles (20 of 454) and 3.1% of paternal uncles (16 of 511) were reported homosexual, a nonsignificant difference, while male cousins ran 4.4% on the maternal side against 5.5% on the paternal side, opposite to the X-linkage prediction. That sample's questionnaire did not separate maternal cousins through aunts from those through uncles, so its cousin contrast is the less precise of the three. In the Gay Pride sample, maternal- and paternal-uncle rates were 2.5% (8 of 320) and 2.1% (6 of 292), and sons of maternal aunts had essentially the same rate as other male cousins, 4.6% (25 of 540) versus 4.4% (49 of 1,120). In the affected-sibling-pair sample, the sample most enriched for whatever familial factor is operating, both comparisons ran opposite to the prediction and by the largest margins observed: maternal uncles 4.5% (5 of 112) against paternal uncles 6.8% (8 of 118), and sons of maternal aunts 4.1% (5 of 122) against other male cousins 7.2% (22 of 307). No contrast in any sample reached significance. The authors correctly noted that realistic X-linked effects of modest magnitude would produce small pedigree differences requiring much larger samples, and that the sibling-pair sample showed signs of ascertainment bias in its unexpectedly high reported rate of homosexual fathers. Their conclusion was that familial factors matter and that X-linked variants, if involved at all, explain only a minority of cases.

The first full-genome linkage scan clarified why ascertainment mattered. Mustanski et al. (2005) analyzed both previously selected families and newly recruited families without a maternal-transmission restriction. The full sample produced its strongest multipoint peaks at 7q36, 8p12, and a maternally influenced region at 10q26; no X-chromosome region exceeded a modest threshold. Yet the previously selected families recovered a strong Xq28 signal when analyzed with the denser original marker set. Marker density and telomeric coverage contributed to the discrepancy, but the reduction after adding unselected families was also consistent with etiological heterogeneity.

A much larger linkage scan later restored support for both an autosomal and an X-linked component. In 409 independent homosexual brother pairs, Sanders et al. (2015) identified primary multipoint peaks in the pericentromeric region of chromosome 8 and at Xq28.

Incorporating liability classes for the fraternal birth-order effect left both peaks essentially unchanged, indicating that the genetic signals were not statistical substitutes for older-brother history. That analytic choice was not incidental. Cantor et al. (2002) had shown that men whose homosexuality is attributable to fraternal birth order constitute phenocopies in linkage analysis, and that because the probability of phenocopy status rises systematically with birth order, the resulting noise is nonrandom rather than random. They proposed exactly the birth-order-dependent liability classes, covariates, and weights that the later scan implemented. The stability of the peaks under that correction is therefore a meaningful result: the chromosome 8 and Xq28 signals are not the immune route wearing a genetic disguise.

These studies justify continued family-based molecular work; they do not establish that advertisement ascertainment has recovered a true Xq28-defined route. That claim remains at risk. A direct test must identify an X-linked molecular signal in pedigrees selected independently of the marker, reproduce it in held-out families, and show that the same profile predicts maternal-line recurrence or sex-antagonistic fitness effects. Otherwise, the selected-family pattern is better attributed to sampling and statistical instability.

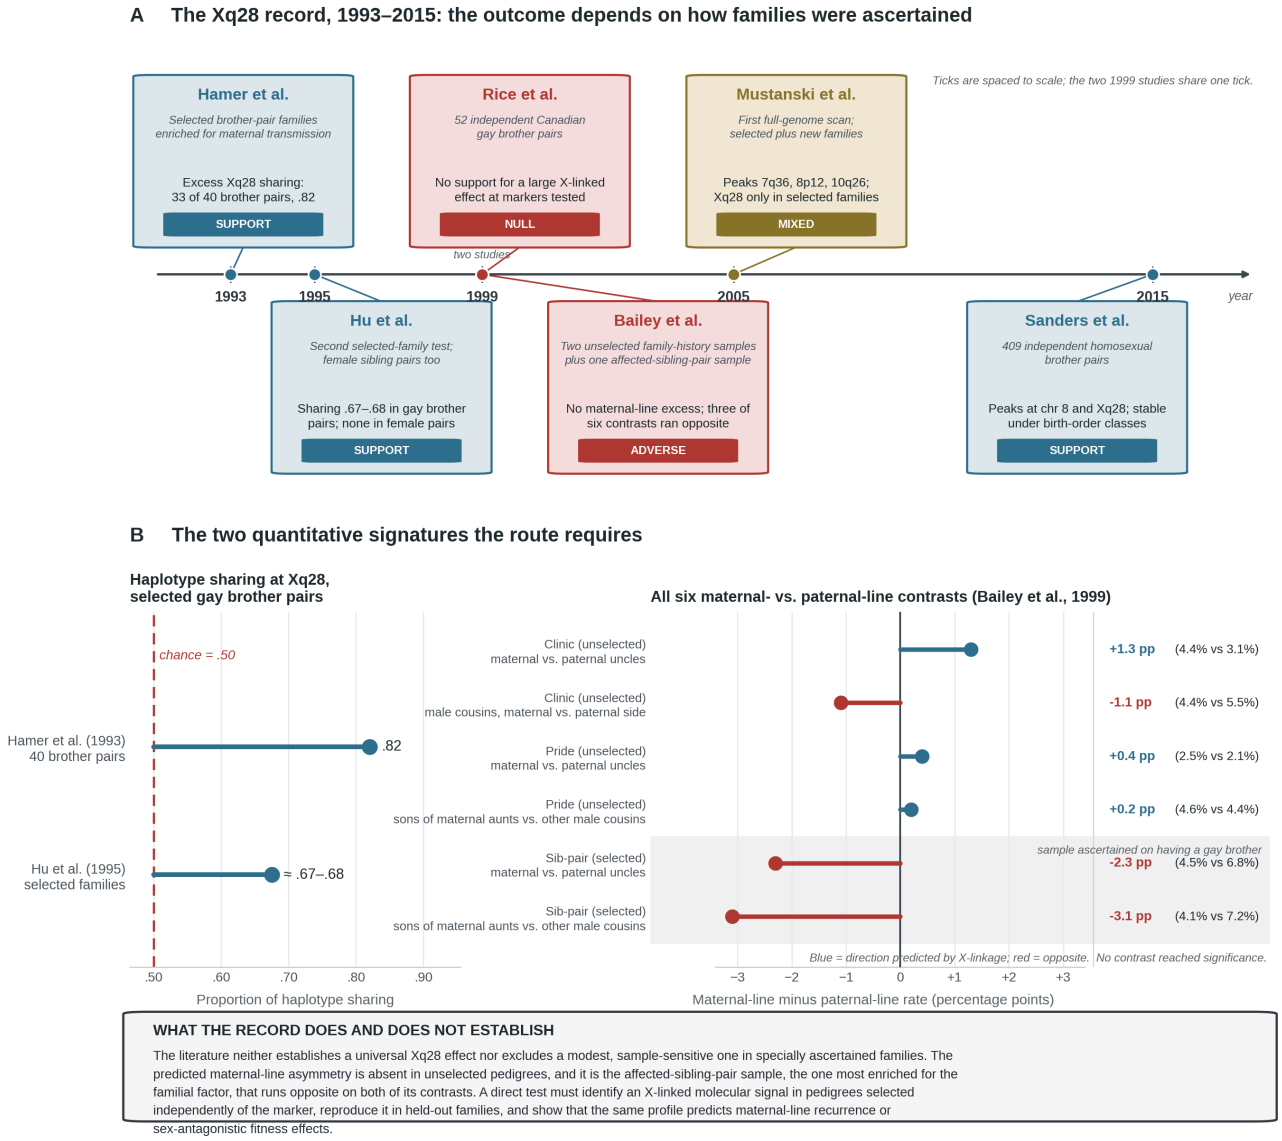


Figure S1 *The Xq28 Linkage Record and Its Dependence on Ascertainment*. Panel A places the principal studies on a timeline whose ticks are spaced to scale; the two 1999 studies share one tick. G. Rice et al. (1999) appears as a null without a plotted estimate because the result is recorded as an absence of support rather than as a point estimate. Bailey et al. (1999) is described as two unselected family-history samples plus one affected-sibling-pair sample, because the third sample was selected on the presence of gay brothers and showed signs of ascertainment bias. Panel B shows the two quantitative signatures an X-influenced route requires. The left plot gives reported haplotype sharing in selected gay brother pairs against the chance expectation of .50. The right plot gives all six maternal-line minus paternal-line differences reported by the source, grouped by sample and labeled for the comparison actually made. The cousin rows are labeled separately because they are not equivalent: in the two interview samples the contrast is sons of maternal aunts versus other male cousins, which is not a maternal-versus-paternal contrast because other male cousins pool maternal-uncle sons with paternal cousins, whereas the clinic questionnaire distinguished only maternal from paternal side. The affected-sibling-pair sample is shaded because it was ascertained on the presence of a gay brother; it is the sample most enriched for the familial factor, and both of its contrasts run opposite to the prediction. No contrast in any sample reached significance.

S4 Non-European Ancestry and Maternal X-Regulatory Variation

The Han Chinese Orientation GWAS and Its Functional Follow-Up

A two-stage genome-wide association study in Han Chinese men used a stringent orientation phenotype, with all cases scoring 5 or 6 on the Kinsey scale and all controls scoring 0. In 1,478 homosexual and 3,313 heterosexual men it reached genome-wide significance at an intronic variant in FMR1NB at Xq27.3 and an intronic variant in ZNF536 at 19q12, and the previously reported Xq28, 7q36, 8p12, and 10q26 regions did not replicate (S.-H. Hu et al., 2021).

The result establishes three limited points. First, a polygenic score trained in a European same-sex-behavior GWAS predicted case-control status in the Han Chinese orientation sample, explaining approximately 1.1% to 2.3% of liability-scale variance; some inherited signal therefore generalizes across ancestry and phenotype definitions. Second, the significant loci differed from earlier linkage regions. That nonreplication can arise from limited power, allele-frequency and linkage-disequilibrium differences, phenotype definition, winner's curse, or etiological mixture; it is not positive evidence for mixture by itself. Third, FMR1NB is at Xq27.3, not Xq28. Its X-chromosomal location does not demonstrate maternal transmission or continuity with the earlier Xq28 linkage claim.

CRISPR-mediated knockout of the entire FMR1NB gene in male mice abolished the significant female partner preference observed in wild-type males, with a nonsignificant trend toward more male-directed mounting in small groups. This is relevant evidence that the gene can influence mammalian mating behavior. It is not functional validation of the causal effect of the associated human intronic variant, because a whole-gene knockout is a much larger perturbation and the causal variant and molecular consequence remain unidentified. The correct inference is mechanistic plausibility at the gene level, not proof that the human association alters partner-sex valuation through the same pathway.

Maternal X-Chromosome Inactivation Skewing

Extreme X-chromosome inactivation skewing, defined as greater than 90% inactivation of one X, was present in 13% of mothers of homosexual sons against 4% of controls, and in 23% of mothers of at least two gay sons; mean inactivation ratios did not differ, so the signal was a heavier tail and greater variance rather than a general shift (Bocklandt et al., 2006). Its limitations are as follows.

The finding is not proof that maternal X-inactivation causes a son's orientation. Peripheral leukocytes are an uncertain proxy for neurodevelopment, some cases came from linkage-enriched families, control offspring were not comprehensively phenotyped, and the extreme cells were small. Its theoretical importance is narrower: maternal X-regulatory variation is concentrated in a familial subgroup and has the distribution expected of a threshold process. It

may mark a segregating X-linked variant, a maternal developmental state, or both. In every interpretation, it supports etiological substructure rather than a universal maternal mechanism.

Relation to Developmental Robustness

Canalization is used here in its conventional sense: developmental robustness to perturbation. A heavier tail of extreme maternal X-inactivation states is compatible with a threshold system, but it does not demonstrate canalization, neural relevance, or transgenerational causation. The result may mark a segregating X-linked variant, a maternal regulatory state, selection in leukocyte lineages, or another process. Its theoretical value is that it identifies unusual regulatory variance in a selected subgroup. The theory requires replication and a demonstrated connection to fetal neural development before treating it as a mechanism-linked indicator.

S5 Population-Genetic Models and Evidence for Carrier Advantage

This section concerns inherited variation. Balancing selection is not a label for maternal immunization, stochastic developmental divergence, or persistence of an immune system. For inherited allele A_j, the selection-only change is $\Delta p_{\text{sel}} = p_j(1 - p_j)\,\bar{s}_j(p_j)$. Balancing requires an internal zero $\bar{s}_j(p_j^*) = 0$ and local stability under the selection-only recurrence. Mutation is added only afterward; an internal equilibrium created while $\bar{s}_j$ is negative at every frequency is mutation-selection balance, not balancing selection.

Miller's Threshold Model

The manuscript adopts the skeleton of an earlier verbal model. Its arithmetic is worth reproducing, because it shows how a variant that eliminates reproduction in a small minority of carriers can nonetheless be maintained, and because its limitations bear on how much weight the present theory can place on it.

The framework has a clear intellectual precursor that deserves explicit acknowledgment. Miller (2000) proposed that sexual orientation is polygenic and developmentally mediated, that the contributing alleles are pleiotropic rather than dedicated, and that in most male carriers they produce moderate shifts toward sensitivity, empathy, tender-mindedness, and reduced aggressiveness. Only when a sufficiently large combination of such alleles co-occurs, potentially together with prenatal or stochastic influences, does development cross a threshold associated with bisexuality or homosexuality. His illustrative five-locus additive threshold model made the arithmetic of balancing selection transparent. Conditional on carrying one focal allele, the 1/16 expression fraction assumes that each of the other four required alleles is present independently with probability .50, so (1/2)^4 = 1/16. Under that illustrative frequency assumption, a reproductive advantage on the order of 16/15 among the far more numerous

heterosexual carriers offsets zero reproduction among the 1/16 in whom the focal allele completes the threshold genotype. Different allele frequencies change the required advantage, often substantially. The equilibrium is a property of the specified example, not a fitted estimate.

Miller's five-locus example is arithmetic, not a fitted population-genetic model. It shows how small benefits among many nonexpressing carriers could offset a large cost in a rare threshold configuration, but it does not establish that the proposed traits are genetically coupled to orientation, that the required benefits existed, or that five additive loci approximate the true architecture. The present theory retains the threshold intuition and places its empirical burden on component-specific genetic and reproductive data.

Formal Models and Distinguishing Predictions

Gavrilets and Rice (2006) developed the essential formal predecessor by comparing overdominance and sexual-antagonism models. The models generate contrasting predictions for chromosomal location, dominance, and effect size of orientation-related loci. Overdominant maintenance is expected to favor different genetic signatures from sex-antagonistic maintenance, and X-linked and autosomal loci experience different sex-specific transmission. These predictions prevent the phrase carrier advantage from becoming a generic explanation: a molecularly defined inherited component must show the genetic properties implied by the proposed equilibrium.

The population-genetic requirements are stricter than a verbal account suggests. Camperio Ciani et al. (2008) compared multiple inheritance models and showed that one-locus autosomal, X-linked, overdominant, and maternal-effect models were generally unable to produce all three required outcomes at once: a stable low-frequency polymorphism, a maternal-line excess of male homosexuality, and higher fecundity in female maternal relatives. Two-locus sexually antagonistic models with one or two X-linked loci did satisfy those conditions across broad parameter ranges. The two-X-locus model provided the closest qualitative match to the observed pedigree pattern. This theoretical result changes the interpretation of maternal-line data. The relevant claim is not “an X-linked homosexuality gene.” It is a multilocus, partially X-linked system whose expression and equilibrium depend on interactions among loci and sex-specific fitness effects.

Social Stratification as a Frequency-Dependent Alternative

Barthes et al. (2013) modeled a stratified society in which female hypergyny can maintain a fertility-associated factor despite a pleiotropic cost expressed as male homosexual preference. The model predicts frequency-dependent maintenance tied to class structure and reported a cross-society association between social stratification and documented male homosexual preference. This is a genuine population-genetic alternative, not a prenatal path. It belongs in the model comparison because it predicts that the sign and magnitude of carrier advantage vary

with social hierarchy and mating structure; failure of a universal female-fecundity effect does not by itself test that environment-dependent mechanism.

Evidence for Carrier Advantage

The empirical support for carrier advantage is real but narrower than usually claimed. In an Australian twin sample, heterosexual men with a nonheterosexual twin reported more opposite-sex partners than controls, and a substantial portion of the covariance between orientation liability and mating success was attributed to shared additive genetic factors (Zietsch et al., 2008). Genome-wide analyses likewise found positive genetic correlations between same-sex behavior and number of opposite-sex partners, although not with number of children in contemporary datasets (Zietsch et al., 2021). Partner number is not reproductive success, especially under modern contraception, but it can preserve a signal of traits that historically increased mating access.

Song and Zhang (2024) demonstrated that the positive fertility association for bisexual behavior is genetically mediated by risk-taking and disappears after controlling that pathway. This is a plausible maintenance mechanism for a behavioral-propensity layer. It does not explain enduring exclusive androphilia. Felesina and Zietsch (2024) proposed that moderate expression of male-attraction liability could enhance paternal-care traits and heterosexual mate value, but the predicted reproductive advantage has not been directly demonstrated.

The current genomic assessment is mixed (Felesina & Zietsch, 2025). Pleiotropic correlations exist, but modern partner number is not historical reproductive success, broad same-sex behavior is not orientation, and correlations do not identify an equilibrium. The theory predicts that the aggregate of selection-valid inherited partitions covering a majority of attraction SNP heritability was nonnegative under natural fertility. The claim is restricted to that covered majority; it does not claim that every associated variant was balanced, that all inherited liability shared one maintenance process, or that covered estimates bound the uncovered fraction.

Coverage, Compensation, and Portability Commitments

The compensatory-maintenance claim is evaluated through mutually exclusive, preregistered partitions of the full common-variant signal, not through a small set of orthogonal components. Partitions may be functional annotations, nonoverlapping linkage-disequilibrium blocks, or disjoint score bins and may contain thousands of correlated small effects. A partition becomes selection-valid only after out-of-sample heritability or score replication, within-family validation, measurement invariance, ancestry-specific linkage-disequilibrium calibration, and transport to natural fertility. Let C be the fraction of out-of-sample SNP heritability covered by selection-valid partitions. The coverage and compensation criteria are those in Table 1 of the main text and are not restated here.

For the covered partitions, K_C is expected reproductive cost attributable to expression in androphilic males and G_C is carrier-weighted gain in all other contexts. Compensation is consequential only if the lower bound on $G_C/K_C \geq 1.0$, unless a separately specified equilibrium demonstrates nonnegative average selection. No selection effect is assigned to the uncovered fraction $1 - C$. The covered set may not be used to manufacture a residual-rescue bound for partitions that failed the validity gates. Sensitivity analyses may display a labeled range of hypothetical uncovered effects, but they do not enter the decision rule.

Portability is a gating condition. Modern scores and partitions must not be transported directly into historical or natural-fertility cohorts without ancestry-specific linkage-disequilibrium recalibration, phenotype measurement invariance, within-family validation, direct genotyping where possible, explicit gene-by-environment and fertility-regime interaction, and modeling of assortative mating. Ancestral imputation is a sensitivity analysis rather than primary genotype. Failure to meet these conditions makes Critical Study 5 inconclusive, not favorable. No existing dataset meets the full measurement, coverage, and natural-fertility requirements. Feasibility planning must simulate discovery samples of 250,000, 500,000, and 1,000,000 or more repeated-attraction phenotypes, plus directly genotyped natural-fertility validation resources, and adopt the lowest design that meets the prespecified precision gates. Critical Study 5 is therefore a long-horizon falsifier.

S6 Maternal-Line Fecundity: Supporting and Contradictory Evidence

Positive Pedigree Evidence

The pedigree studies described in the manuscript found higher fecundity among mothers and maternal aunts of homosexual men. The strongest version of that test examines relatives who were not selected by virtue of having borne a homosexual son.

Camperio Ciani and Pellizzari (2012) supplied that version by examining completed fecundity in nonparental female relatives. Maternal aunts and maternal grandmothers of homosexual men had more children than their counterparts in heterosexual pedigrees, while paternal aunts and grandmothers did not. Within families, the direction of the maternal-paternal contrast reversed between homosexual and heterosexual probands in the manner predicted by the X-linked sexually antagonistic models.

Direct Adverse Evidence

The negative record is more damaging than a generic statement that some studies failed to replicate the effect. Blanchard et al. (2020) applied one analytic method to archived data from 10 studies comprising 14 samples and 5,390 participants. The pooled analysis supported the fraternal-birth-order effect but did not support the female-fecundity effect; the two proposed

mechanisms separated exactly where the theory says they should be separable. Semenyna et al. (2023) then jointly re-examined the two effects across a decade of Samoan data. The older-brother association was retained, whereas the female-fecundity effect again was not supported.

Other adverse evidence points in the same direction. Fořt et al. (2024) reported equal or lower fertility among mothers of gay men after accounting for birth-order effects. Ablaza et al. (2022), using Dutch administrative data on more than nine million individuals, found that larger sibships predicted lower rather than higher odds of entering a same-sex union; their models supported independent later-birth-order and older-brother effects but rejected a simple fecundity inference from proband sibship size. Bailey et al. (1999) found no maternal-line excess in three independently ascertained samples, with two comparisons running opposite the prediction.

The current pooled evidence is therefore unfavorable to a general female-fecundity mechanism. A rare, molecularly defined maternal-line inherited component remains logically possible because population and proband-based samples can dilute a concentrated signal, but no study has yet shown that selecting families on a direct X-linked or sex-antagonistic molecular marker rescues the fecundity effect. Until that is demonstrated, increased maternal-line fecundity is a high-risk prediction of one proposed inherited component, not an established pillar of the portfolio.

S7 Kin-Directed Investment: Positive and Negative Evidence

The avuncularity hypothesis has been tested most carefully in Samoa, but a direct negative study in a Western community sample and the inclusive-fitness arithmetic constrain its role. It is treated as a subsidiary, culture-contingent channel rather than a general explanation of persistence.

The proposal holds that androphilic males offset reduced direct reproduction by enhancing the reproductive success of close relatives. The best empirical test comes from Samoa, where fa'afafine constitute a socially recognized, overwhelmingly androphilic category of biological males. Vasey and VanderLaan (2010) replicated in an independent sample of 56 fa'afafine and 95 gynephilic men the finding that fa'afafine report greater willingness to invest in nieces and nephews, $M = 6.27$, $SD = 1.04$, versus $M = 5.62$, $SD = 1.30$, $t(135.66) = 3.49$, $p = .001$, Cohen's $d = .57$. Critically, they then addressed the obvious alternative explanation. Combining samples to yield 94 fa'afafine, 66 childless gynephilic men, and 72 gynephilic fathers, and covarying age and income, the group effect remained significant, $F(2, 227) = 9.18$, $p < .001$, with fa'afafine exceeding both childless gynephilic men, $d = .65$, and gynephilic fathers, $d = .56$, while the two gynephilic groups were indistinguishable, $d = .01$. Among gynephilic men, number of children was unrelated to avuncular tendencies, $r = .02$, $p = .81$. Elevated avuncularity among fa'afafine is therefore not a simple consequence of being unencumbered by children of one's own.

Bobrow and Bailey (2001) obtained the opposite result in a community sample of homosexual and heterosexual men. Homosexual men were no more likely to channel resources toward relatives; heterosexual men tended to give more financial resources to siblings, and homosexual men reported greater estrangement from some family members. The study did not measure recipient fertility and was conducted in a social environment very different from Samoa, but it demonstrates that elevated kin investment is not a general consequence of male homosexuality.

The inclusive-fitness arithmetic nonetheless constrains how much this mechanism can carry. An uncle shares approximately one quarter of his genes with a niece or nephew, so a nonreproducing male must facilitate roughly two additional nieces or nephews for every forgone child before the ledger balances. The observed effects were moderate, the measure was stated willingness rather than resources actually transferred, and recipient reproductive outcomes were not quantified. Vasey and VanderLaan were explicit that consistency with a kin-selection prediction does not demonstrate that androphilia evolved as an adaptation for kin-directed altruism, and they noted the alternative that elevated avuncularity amplifies the effects of genetic factors that increase female-relative fecundity without itself being the selected function.

A credible maintenance test must therefore measure actual time and material transfers, compare androphilic men with childless gynephilic men and women, and estimate the effect on recipient survival and reproduction. Stated willingness and family sentiment are insufficient. Even a replicable group difference would remain evolutionarily negligible if it failed to approach the inclusive-fitness threshold after cultural variation, estrangement, opportunity costs, and uncertainty in counterfactual reproduction are incorporated.

S8 Childhood Gender Nonconformity: The Bivariate Twin Evidence

Whether childhood gender nonconformity and adult orientation share causes, rather than merely co-occurring, is answered by a bivariate genetic design. The estimates, the corroborating evidence for their nonuniversality, and their limitations are given here.

The decisive question is not whether the association exists but whether the two phenotypes share causes. Alanko et al. (2010) answered it with a bivariate twin design in a population-based Finnish sample. Twin correlations for recalled gender-atypical behavior were .56 in monozygotic and .27 in dizygotic male pairs; for adult orientation they were .50 and .25. Across parameterizations, genetic effects accounted for roughly 29% to 51% of variance in each phenotype, shared-environment estimates were small or zero, and roughly half or more of the variance fell to nonshared environment inclusive of measurement error. In the best-fitting male model, an AE model, the genetic correlation between the recalled childhood measure and adult same-sex orientation was +.73 on the girlishness scoring described in S1, while the nonshared

environmental correlation of .15 was nonsignificant. For women the best-fitting model was a DE model with a dominance correlation of −.30 and a small but significant nonshared environmental correlation of −.12; rescored to the atypicality direction for each sex, both genetic correlations are positive and the sex difference is one of magnitude.

The fitted male genetic correlation of .73 does not mean that 73% of genetic variance is shared. Under the model, its square is approximately .53, so a little over half of the standardized genetic variance was common to the two measured phenotypes. The estimate came from one sample with retrospective measurement, a 36% response rate, questionnaire-based zygosity, past-year rather than lifetime orientation, and nonshared environment confounded with measurement error. It is strong evidence of partial common inherited influence under that model, not a universal coefficient or proof of a prenatal mechanism.

The nonuniversality is corroborated across designs. Orientation accounted for only a minority of variation in childhood sex-typed behavior, with substantial overlap within every group (Bailey & Zucker, 1995). In the Australian twin register, childhood gender nonconformity was familial and heritable and covaried with orientation, but individual-specific influence remained large (Bailey et al., 2000). In Bailey and Pillard (1991), a proband's childhood gender nonconformity did not predict homosexuality in his relatives even though nonconformity was highly correlated within concordant monozygotic pairs.

Childhood gender nonconformity is therefore an early correlate and possible developmental trace. It is neither a compulsory precursor nor an etiological class indicator by itself. Its value in the causal graph is multivariate: when combined with direct familial, endocrine, immune, and regulatory measures, it can help predict external outcomes and distinguish path loadings; alone, it cannot assign cause.

S9 The Fraternal Birth-Order Association: Quantification and Alternatives

Population Attributable Fractions

Two published estimates of the share of male homosexuality attributable to fraternal birth order differ by roughly a factor of two. Their derivations are given here, because the manuscript's claim that these fractions are population-level quantities that assign no individual case depends on how each was computed.

Under their causal and log-linear assumptions, the two published birth-order models imply a minority population contribution and bracket its magnitude. Using two national probability samples, Blanchard and Bogaert (2004) found homosexuality rates of 2.18% among men with no older brothers and 4.83% among men with at least one, an odds ratio of 2.28, and estimated a population attributable fraction of 28.6%, with a wide 95% interval from 14.8% to 48.0%. Cantor

et al. (2002) approached the same quantity with a method designed for an ordered exposure rather than a binary one. Generalizing the epidemiological attributable fraction to a variable with multiple exposure levels, and assuming a log-linear relationship between number of older brothers and the odds of homosexuality, they derived the attributable fraction at exposure level n as a function of baseline prevalence and the per-brother odds ratio. In 302 homosexual and 302 birth-year-matched heterosexual men, with an assumed baseline prevalence of 2% and a logistic per-brother odds ratio of 1.33, the attributable fraction rose from 0 with no older brothers to .243 with one, .426 with two, .563 with three, .667 with four, and .745 with five or more. Weighting by the observed distribution of older brothers among the homosexual participants yielded a population attributable fraction of 15.1%. The estimate was robust to the baseline assumption, changing only to 15.2% at 1% prevalence and 14.8% at 4%, and the log-linearity assumption survived direct testing: after the linear term, neither a quadratic, $\chi^2(1) = 0.24$, $p = .62$, nor a cubic term, $\chi^2(1) = 0.23$, $p = .63$, improved prediction.

Cantor et al. (2002) also defined AF50, the older-brother exposure at which the model-based attributable fraction among exposed cases reaches .50; their estimate was approximately 2.503 older brothers. This is a population-level quantity under the fitted causal and log-linear assumptions. It is not an individual probability that a particular man's orientation was caused by an immune mechanism. Under the fitted causal model, older-brother exposure accounts for a larger fraction of outcome risk in high-exposure strata; it does not identify the mechanism or assign individual causation from brother count alone.

Alternatives, Statistical Critiques, and Immune Evidence

The manuscript treats observed birth-order coefficients as mixtures rather than as pure markers of maternal immunization. The competing account, the register evidence, and the statistical critiques are set out here, along with the evidence that nonetheless is consistent with a pregnancy-history-dependent immune path.

An immunological interpretation of birth order is not the only one available, and the alternative deserves a hearing precisely because it is testable. Miller (2000) accepted the birth-order association but proposed that increased androphilia among later-born sons is a by-product of a maternal mechanism for diversifying offspring phenotypes. On this account, later-born sons are shifted toward greater flexibility, cooperation, sensitivity, and lower dominance, which reduces direct competition with older brothers for the same familial and social niches; homosexuality is the tail of that distribution rather than the target of the mechanism. The two hypotheses make divergent predictions that no published study has adjudicated. The immune account predicts that birth-order effects should be specific to male fetuses, should track maternal antibody titers, and should not extend to sex-dimorphic personality traits in heterosexual later-born sons. The diversification account predicts precisely such extension, and predicts it in the absence of antibody elevation.

The birth-order evidence is in fact more complex than a strictly male-specific immune account predicts. Ablaza et al. (2022) isolated older-brother, general birth-order, and sibship-size effects in Dutch population registers. They found independent older-brother and later-birth-order associations with same-sex union, but the older-brother coefficient was similar in men and women. The comparable effect in women challenges a model in which the entire population association is caused solely by antibodies to male-specific antigens acting on male fetal brains. It implies either that the register phenotype captures additional mechanisms, that maternal immune responses can spread to non-sex-specific targets, or that some birth-order effects arise through pathways not represented in the classic hypothesis, including the sibling-diversification pathway.

Methodological critiques further show that standard controls for family size can create or exaggerate birth-order associations under some conditions (Vilsmeier et al., 2023). The correct conclusion is neither that the fraternal birth-order effect is universally decisive nor that it is an artifact. The biological-brother specificity in many studies, the NLGN4Y evidence, and its independence from linkage signals under liability-class correction are consistent with a maternal immune path. The Dutch register findings and statistical critiques demonstrate that older-brother count is not a pure biomarker of that path in every dataset.

Blanchard, Kabátek, and Zdaniuk (2025) provide the current ABZ logistic procedure for disentangling older-brother, older-sister, and sibship-size associations. Its ceteris-paribus coefficients correspond to explicit substitutions within a fixed sibship rather than to vague additions of one sibling, and the authors recommend theory-neutral labels because a sibship-size coefficient does not itself demonstrate female fecundity. The method improves statistical separation of observed associations; it does not convert any coefficient into a maternal-immune mechanism or an individual probability of causation.

Bogaert et al. (2018) reported greater maternal immune responsivity to selected NLGN4Y isoform sequences in mothers of homosexual sons, particularly those with older brothers. The supporting information was corrected in 2022 to restrict the antibody claims to the appropriate sequences (Bogaert et al., 2022). The study was cross-sectional and used current maternal samples to infer an immune state during pregnancies decades earlier. Current titers can be exploratory familial markers; they cannot reconstruct the concentration, timing, tissue access, or specificity of exposure during the index pregnancy.

Failure of NLGN4Y mediation would reject that specified antigen model, not every maternal immune mechanism. Other Y-linked antigens, cell-mediated immunity, fetal microchimerism, pregnancy loss, nonimmune physiological changes, parity structure, and residual confounding remain alternatives. Conversely, replication of an older-brother coefficient in the absence of the predicted antibody pattern would not positively identify Miller's sibling-diversification account. The prospective serial-sampling design in Section S13 distinguishes these inferential levels.

S10 The Postmortem and Neuroimaging Record

Postmortem Neuroanatomy

Three findings anchor the postmortem literature: LeVay (1991) reported a smaller third interstitial nucleus of the anterior hypothalamus in homosexual men; Allen and Gorski (1992) reported a larger midsagittal anterior commissure; and Swaab and Hofman (1990) reported an enlarged suprachiasmatic nucleus. S.-H. Hu et al. (2021) did not replicate the original SCN volumetric finding. They reported lower average ZNF536 immunoreactivity within the SCN of homosexual donors. The molecular observation and the earlier volume claim must remain distinct.

Each of these studies shares a set of limitations that are rarely stated together. Allen and Gorski (1992) assembled 30 age-matched triplets from an initial pool of 256 postmortem specimens, and their measurements were substantial: mean anterior commissure area was 14.20 mm^2 in gay men, 10.61 mm^2 in heterosexual men, and 12.03 mm^2 in heterosexual women, an omnibus effect at $p = .0001$, with the gay male mean 34% larger than the heterosexual male mean. Tracing reliability was high, $r = .94$. But 24 of the 30 gay men had died with AIDS, against 6 of the 30 heterosexual men; sexual orientation was inferred from medical records, and in most comparison subjects heterosexuality was assumed from the absence of documentation to the contrary; there was no lesbian group; and the midsagittal plane was located by hand. The authors argued that AIDS was unlikely to explain the enlargement, noting that AIDS neuropathology more typically produces atrophy and that limited histological examination found no edema or vacuolization. That argument is reasonable and remains unconfirmed.

The most rigorous attempt at replication produced a partial reversal that deserves far more prominence than it receives. Byne et al. (2001) examined all four interstitial nuclei in 34 presumed heterosexual men, 34 presumed heterosexual women, and 14 homosexual men, and went beyond volume to count neurons. They confirmed a robust sex difference in INAH3: mean volume was 0.121 mm^3 in men and 0.073 mm^3 in women, $F(1, 61) = 20.7$, $p < .0001$, and, more informatively, INAH3 contained roughly 60% more neurons in men than in women, with mean estimates of 1,794 and 1,123. INAH1, INAH2, and INAH4 showed no significant volumetric sex difference. For sexual orientation, however, the result was null. Homosexual men had an intermediate mean INAH3 volume of 0.096 mm^3, but neither the volume nor the neuron count differed significantly from heterosexual men in pairwise comparison; the mean neuronal estimate in homosexual men, 1,831, was numerically the highest of the three groups. Only neuronal packing density showed a near-significant group effect, $F(2, 53) = 3.03$, $p = .057$. The authors concluded that any orientation-associated difference might involve neuropil or extracellular spacing rather than neuron number, and stated explicitly that INAH3 function was unknown and that orientation could not be reliably predicted from nuclear volume.

Byne et al. also could not separate orientation from disease: all 14 homosexual cases were HIV-positive and had died of opportunistic infections, because orientation was classified from documented HIV risk factor, with intravenous drug use coded as presumed heterosexual and homosexual behavior coded as homosexual. No women of known homosexual orientation were available. The same confound structure recurs in the recent Netherlands series, in which all 13 homosexual donors and 4 of 13 heterosexual donors had AIDS (S.-H. Hu et al., 2021).

The evidentiary summary is therefore narrow. INAH3 is sexually dimorphic, but an orientation difference in volume or neuron number was not confirmed in the best-controlled replication. The anterior commissure difference was reported once and remains unreplicated. The SCN volume difference was reported once; a later study found a different molecular feature in the same structure. Nearly all homosexual postmortem cases in these series were confounded with AIDS or HIV-related ascertainment. No postmortem marker has been shown to classify orientation in a representative, disease-unconfounded sample.

Human Neuroimaging and Distributed Systems

Savic and colleagues reported orientation-associated differences in responses to putative chemosensory cues, cerebral asymmetry, and amygdala connectivity (Savic et al., 2005; Savic & Lindström, 2008). Frigerio et al. (2021) reviewed structural, functional, and metabolic imaging studies more broadly. The literature contained conflicting results, considerable overlap between homosexual and same-sex heterosexual groups, small and unrepresentative samples, and no consistently distinctive brain phenotype. That review supplies the correct evidentiary frame: reported differences are hypotheses about distributed systems, not a settled neurobiological signature.

Abé et al. (2021) analyzed magnetic-resonance imaging and genotype data from 18,645 UK Biobank participants. Multivariate sex-classifying patterns were less pronounced among participants reporting same-sex behavior, especially women, and same-sex-behavior polygenic scores were associated with lateral occipital and temporo-occipital volume. The sample size is a major advance, but the phenotype was broad same-sex behavior, the effects were small, and the design was cross-sectional. It supplies a large-scale brain-genomic association, not a route-specific neural signature or evidence that one cross-sex shift causes enduring male androphilia.

Pair-bonding research implicates interacting oxytocin, vasopressin, dopamine, hypothalamic, striatal, and social-reward systems (Blumenthal & Young, 2023). These networks are plausible endpoints through which prenatal androgen signaling, immune perturbation, and polygenic neurodevelopmental variation could change partner-sex valuation. Plausibility is not localization. The theory predicts that stratification by independently measured upstream mechanisms will improve reproducibility; it also accepts the possibility that no path produces a large gross-anatomical marker.

S11 Causal Paths, Prenatal Organization, and Model Comparison

Intellectual Predecessors and the Additional Burden of the Present Theory

Rahman and Wilson (2003) offered an earlier broad psychobiological synthesis spanning genetics, prenatal hormones, neurodevelopment, and evolutionary accounts. VanderLaan et al. (2022) then argued explicitly that similar same-sex orientations may arise through different genetic, hormonal, and immunological pathways and recommended subgroup-oriented research. Swift-Gallant et al. (2019) and Folkierska-Żukowska and Dragan (2025) showed that profile analyses using familiality, birth order, handedness, gender expression, and related proxies are feasible. Causal subgrouping is therefore an extension of an established program, not a novel premise.

The causal-portfolio theory adds four commitments that must be empirically consequential or the theory contributes only relabeling. First, sources, prenatal mechanisms, developmental intermediates, latent attraction, and observed phenotypes occupy different causal levels. Second, a route is a directed path through those levels and requires a numerical material effect, prospective mediation robust to prespecified mediator-outcome confounding, and at least one identification source independent of sequential ignorability; mediation alone is not enough. Third, prevalence is a joint-distribution expectation and route attribution is counterfactual rather than a natural case count. Fourth, inherited and recurrent developmental components are paired with different evolutionary accounting rules.

The Five Causal Levels

The five levels are: (1) sources of variation, including offspring inherited genotype, maternal genotype and physiology, pregnancy history and exposure, and stochastic inputs; (2) prenatal mechanisms, including steroid signaling, placental transport and metabolism, maternal immune activity, and gene regulation; (3) developmental intermediates, including sex discrimination, male- and female-cue incentive salience, attachment, and social reward; (4) latent attraction, represented by repeated male-directed and female-directed attraction; and (5) observed phenotypes and measurement rules, including fantasy, behavior, identity, partnership, and categorical classification. Latent attraction is not folded into its indicators because the phenotype-production and measurement models depend on that distinction.

This separation changes the interpretation of the previous route labels. Diffuse polygenicity is a source architecture that can feed several mechanisms. Maternal-line concentration is a pedigree signature potentially produced by X linkage, maternal genotype, imprinting, autosomal interaction, correlated exposure, or ascertainment. Maternal immune activity is a prenatal mechanism linked to pregnancy history. Regulatory state is a mechanism or modifier. None is a parallel circle at the same causal level.

Prenatal Endocrine Organization and Male Heterogeneity

Swift-Gallant et al. (2023) review experiments of nature, endocrine conditions, and retrospective biomarkers and conclude that early gonadal hormone action influences human sexual orientation. Evidence is particularly strong that androgens increase gynephilia in people raised as girls, and there is suggestive evidence that early estrogens increase androphilia. Adult hormones regulate libido but do not reliably reverse partner preference. These findings strengthen prenatal organization as the central causal stage.

The same review shows why male variation cannot be reduced to one monotonic androgen scale. Evidence among anatomically typical males is mixed, biomarkers are indirect and tissue-insensitive, and multiple male pathways are plausible. Tissue-specific receptor sensitivity, developmental timing, immune perturbation, placental biology, regulatory state, and potentially nonlinear low- and high-androgen effects can generate distinct configurations. The theory predicts configurational organization of cue valuation and social reward, not global under-masculinization.

Why Distributional Class Recovery Is Not Causal Identification

Nonnormal continuous distributions can generate apparently optimal latent classes, and class-restricted analyses can discretize ordinary dimensions (Bauer & Curran, 2003). The reverse error is equally serious: multiple causal paths can converge onto one continuous outcome distribution. A class solution does not prove routes, and a continuous-factor solution does not prove one etiology.

The decisive comparison is among structural causal models: route-free models with exactly one or two latent continuous liabilities; three- to five-liability sensitivity models; broad polygenicity plus measurement differences; independent genetic and gestational mechanisms without an integrated portfolio; and the causal-portfolio graph. Latent classes, factor mixtures, and grade-of-membership models remain useful descriptions but must be embedded in the causal comparison. The selected procedure must recover known simulated structures under nonnormality and measurement error, replicate in probability-weighted hold-outs and across genetic-similarity strata, and predict outcomes not used in estimation.

Discriminating Predictions

The causal-portfolio theory requires at least two mechanism-distinct prenatal paths that make material contributions to repeated attraction and produce different temporal, conditional-independence, external-prediction, or identified edge-intervention patterns. Material is defined as a standardized mediated or identified edge effect of at least .10, with a lower 90% confidence bound above .05 in an independent hold-out. The effect must have a robustness value of at least .20 in the prespecified sensitivity analysis and receive support from a within-family or within-mother contrast, valid quasi-experiment or instrument, targeted perturbation, or prespecified

negative-control restriction. Effects below .03 are practically null; .03 to .10 is indeterminate. A route-free one- or two-liability model permits dimensional structure but no source-specific mediator map. A broad polygenic-plus-measurement model predicts that direct endocrine, immune, placental, and regulatory measures add no qualifying external prediction once reliable attraction and polygenic liability are modeled. Independent-mechanism models predict main effects without requiring calibrated interactions or joint intervention accounting. Mutation-selection models predict negative inherited selection without compensatory carrier benefit. These alternatives are implemented in Critical Study 1.

S12 Comparative Mate-Choice Theory and Speculative Social Extensions

Same-Sex Behavior Is Not Same-Sex Attraction

Lerch (2026) connects the evolution of same-sex sexual behavior to classic mate-choice theory. Sex discrimination is itself an evolved preference, and imperfect discrimination can be favored when choosiness is costly, encounters are rare, mating effort is high, sex signals overlap, or the opportunity cost of rejecting a potential mate exceeds the cost of an unproductive attempt. Acceptance-threshold and signal-detection models therefore predict same-sex behavior without requiring same-sex attraction.

Lerch explicitly states that the review does not provide a direct account of human sexuality and devotes little attention to same-sex attraction. Its relevance to the present theory is nevertheless substantial. It supplies a formal behavioral path from general mating strategy, sex discrimination, and social context to same-sex behavior that bypasses enduring male-directed valuation. It therefore reinforces the requirement to place behavior downstream from attraction and general behavioral propensity rather than use behavior as the phenotype-production variable itself.

The Cayo Santiago Macaque Study

Clive et al. (2023) observed 236 free-ranging male rhesus macaques and found same-sex mounting to be common, repeatable, modestly heritable, and not detectably associated with reduced reproduction in that sample. Same-sex mounting frequency between pairs was associated with coalition frequency after adjustment for relatedness, rank, and age. The study demonstrates that same-sex behavior can be recurrent and heritable without an observable reproductive cost in one population. It does not show that mounting causes coalitions, that the behavior is an adaptation, or that macaques possess an analogue of enduring human orientation.

Male-Oriented Rams

Approximately 8% of rams in studied populations show an enduring preference for other males in standardized choice tests even when estrous females are available (Roselli et al., 2011). Neuroendocrine and preoptic-area differences are consistent with prenatal organization, although the causal sequence remains unresolved. The ram model is the strongest comparative demonstration that stable male-oriented partner preference can arise in a mammal without human identity or culture. It supports the plausibility of the latent attraction endpoint while leaving homology and evolutionary maintenance open.

TRPC2, Sex Discrimination, and Phylogenetic Caution

TRPC2 disruption in mice increases same-sex mounting while conditioned-aversion evidence indicates that sex discrimination can remain intact (Pfau et al., 2021a). Catarrhine TRPC2 pseudogenization is therefore a plausible permissive event affecting behavioral inhibition or developmental organization. Same-sex behavior also occurs in taxa with intact TRPC2, and adult vomeronasal removal does not reproduce every developmental effect (Clive et al., 2021; Pfau et al., 2021b). The hypothesis is a subsidiary sensory-pathway prediction, not a maintenance mechanism for human androphilia.

Adaptive and Social Same-Sex Behavior

Lerch (2026) reviews cases in which same-sex behavior may alter aggression, reconciliation, coalition function, mate-choice copying, sperm transfer, or biparental care. These examples show that same-sex behavior can be a byproduct of an adaptive mating rule or can itself have context-specific benefits. They do not establish a general fitness benefit to enduring male-directed attraction. The act, the attraction, and the inherited liability must remain separate units of analysis.

Reciprocal Alliance Accounts

Kirkpatrick (2000) argued that homosexual behavior may maintain same-sex alliances with reproductive benefits and found more anthropological support for reciprocal-altruism than for simple kin selection. The hypothesis concerns behavior, affiliation, and alliance function across culturally variable settings; it does not by itself explain stable male-directed attraction or identify inherited variants. It enters the portfolio only if alliance-linked behavior yields measured reproductive returns and those returns covary with the inherited liability that incurs the orientation cost.

Indirect Genetic Effects and Kin Investment

Behavior generated by one individual can alter the social and reproductive environment of others, creating indirect genetic effects (N. W. Bailey & Zuk, 2009). No sexuality-related indirect genetic effect has been estimated in any species. Coalition associations, altered mating markets,

alloparental labor, and social tolerance are therefore excluded from the central maintenance claim until pedigreed or genomic network designs show that one individual's genotype or path score changes another individual's completed reproduction.

The kin-investment hypothesis has mixed evidence. Samoan fa'afafine report greater avuncular willingness than comparison men, but a Western community study did not find greater family-directed resource transfer, and neither literature establishes the approximately two additional nieces or nephews per forgone child required by the simplest inclusive-fitness arithmetic (Bobrow & Bailey, 2001; Vasey & VanderLaan, 2010). A valid test measures actual transfers and recipient reproductive outcomes in several kin-dense societies.

S13 Formal Causal-Portfolio Accounting and Full Critical-Study Protocols

The formalization separates causal phenotype production, measurement, prevalence, inherited allele dynamics, recurrent developmental incidence, node interventions, and edge-specific path effects. The equations are not evidence that the parameters exist. They state which quantities must be identified, when identification fails, and which adverse patterns reject the theory.

S13.1 Structural Causal Model

For male i, let S_i contain offspring inherited genotype G_i, maternal genotype and physiology MG_i, pregnancy history and exposure H_i, and stochastic developmental inputs Q_i. Let M_i contain prenatal steroid signaling, placental transport and metabolism, maternal immune activity, and gene regulation. Let D_i contain sex discrimination, male-cue incentive salience, female-cue incentive salience, attachment, and social reward. Let $A_i = (A_{Mi}, A_{Fi})$ be latent attraction. A minimal acyclic representation is:

$$M_i = f_{\mathrm{M}}(S_i, U_{Mi}); \; D_i = f_{\mathrm{D}}(M_i, S_i, U_{Di}); \; A_i = f_{\mathrm{A}}(D_i, M_i, S_i, U_{Ai}).$$

A route p is a directed edge set from one element of S_i through one or more elements of M_i and D_i to A_i. A route is distinct only if it implies a nonredundant temporal, conditional-independence, external-prediction, or identified edge-intervention pattern. Prospective mediation is required but not sufficient: it must survive the prespecified mediator-outcome confounding sensitivity threshold and be supported by an identification design independent of sequential ignorability. Different sources can share mechanisms, different mechanisms can converge, and simultaneous path loadings and nonlinear interaction are permitted.

S13.2 Attraction, Behavior, and Observed Classification

Let $A_{Mi}(t)$ and $A_{Fi}(t)$ denote latent male-directed and female-directed attraction at measurement occasion t, estimated from multiple reliable indicators. Longitudinal summaries $\bar{A}_{Mi}$ and $\bar{A}_{Fi}$ are the primary phenotypes. An observed category Y_i is generated by a prespecified rule m applied to the repeated profile:

$Y_i = m(\{A_{Mi}(t), A_{Fi}(t)\}$ for $t = 1, \ldots, T)$.

Same-sex behavior is downstream. Let L_i index general mating effort and behavioral propensity, V_i opportunity and partner availability, C_i social or legal constraint, and U_{Bi} unmeasured behavioral inputs:

$B_i = g(A_{Mi}, A_{Fi}, L_i, V_i, C_i, U_{Bi})$.

This model makes the distinction between attraction and behavior causal rather than semantic. A one-factor genetic model may explain covariance among A_{M}, A_{F}, and B while the observed phenotypes remain noninterchangeable. Empirical Proposition P2 is tested separately in S13.8.

S13.3 Prevalence, Node Interventions, and Path-Specific Effects

Let Z contain the measured sources, mechanisms, intermediates, covariates, and interactions, and let F_t be their joint distribution in cohort t. Observed prevalence is:

$\pi_t = E_Z[\Pr(Y = 1 \mid Z)]$, the expectation taken over the cohort-t joint distribution F_t.

This is descriptive marginalization over the joint distribution, not a controlled-direct-effect estimand, and defining Z to include intermediates does not license causal conditioning on them. The same prevalence can be produced by different joint distributions, and continuous interaction prevents unique route-specific case counts. Binary sufficient-cause events can be combined by inclusion-exclusion, but that representation is an optional special case and does not identify continuous path incidence.

For a prespecified source or mechanism X and reference x_0, the node-level prevalence effect is $\Delta_{\mathrm{node}}(X, t) = \pi_t - \pi_t\{\mathrm{do}(X = x_0)\}$. This is a total effect of intervening on X, not a path-specific effect (Pearl, 2009; Hernán & Robins, 2020). Under this sign convention, a protective reference intervention that lowers prevalence yields a positive Δ_{node}.

For a prespecified directed edge set P, define:

$\mathrm{PSE}_{P,t} = E_t[\bar{A}_{\mathrm{M}}(x \text{ along } P, x_0 \text{ along } \bar{P}) - \bar{A}_{\mathrm{M}}(x_0)]$.

The primary edge-intervention contrast is defined on latent attraction; the corresponding contrast on Y is a derived quantity obtained by applying m. The node, edge set, mediator order, target population, and allocation of interactions must be preregistered. Natural and path-specific effects require assumptions stronger than temporal ordering. Recanting witnesses can prevent identification of edge-specific effects, and an exposure-induced mediator-outcome confounder can prevent ordinary natural-effect decomposition even when measured (Avin et al., 2005; Robins & Greenland, 1992; VanderWeele et al., 2014). Every candidate mediated effect is summarized by its robustness value, the smallest assumed absolute residual mediator-outcome correlation at which its lower 90% bound falls to .05, and fails the gate if that value is below .20 (Cinelli & Hazlett, 2020; VanderWeele & Ding, 2017). The full curve is reported rather than a single pass-or-fail verdict. Because that analysis does not prove ignorability, at least one within-family or within-mother contrast, valid quasi-experiment or instrument, targeted perturbation, or prespecified negative-control restriction must independently support the path. When identification fails, report node or joint-intervention effects, interventional analogues, or bounds rather than path-attributable case fractions (Imai et al., 2010; VanderWeele, 2015). A Shapley allocation is a convention-dependent descriptive partition, not a discovered causal fraction (Shapley, 1953).

Worked SCM: Same Data, Different Identification Assumptions

Simulation recipe. Set seed 20260825 and n = 20,000. Draw independent standardized sources G and I. Generate $M_G = .40G + e_G$ and $M_I = .40I + e_I$, with residual variances chosen to standardize both mechanisms. Generate $A = .25M_G + .25M_I + e_A$, again standardized. Each true source-to-mechanism-to-attraction indirect effect is .40 × .25 = .10. Generate four attraction indicators with reliability .70, split the one dataset 70:30, estimate every model in the training sample, and evaluate all models in the same hold-out.

Same-data model comparison. The strict portfolio SCM yielded held-out Gaussian log likelihood −9.870 per observation. The strict route-free SCM, in which G and I predicted A directly and M_G and M_I had no remaining association with A, yielded −9.922. The strict portfolio advantage was therefore .052 log-likelihood units per observation. A route-free sensitivity SCM that retained the direct source effects but estimated residual M_G–A and M_I–A correlations yielded −9.870, essentially tying the portfolio model; the fitted residual correlations were .225 and .236. The strict route-free model is distinguishable, but the confounded route-free model is observationally competitive.

Confounding degradation. In a second simulation, the structural coefficients remained .40 and .25, so each causal indirect effect remained .10, but the residual mediator-outcome correlation was set to .10. A naive no-confounding mediation analysis estimated the two indirect effects at approximately .136 and .136. At residual correlation .20, the corresponding estimates were approximately .167 and .168. These are not small distortions around the decision boundary. The

example therefore demonstrates both facts the theory must respect: explicit route restrictions can defeat a strict route-free model, and modest unmeasured *M–A* confounding can reproduce or amplify the apparent portfolio signature. Critical Study 1 accordingly requires confounding-robust mediation plus independent identification.

S13.4 Inherited Allele Dynamics

For inherited component A_j with frequency $p_{j,t}$, let $W_{Aj,t}$ be expected reproductive fitness of carriers averaged across sex, expression, path background, and demographic regime. The deterministic selection recurrence is:

$$p_{j,t+1} = p_{j,t}\, W_{Aj,t} / \bar{W}_t.$$

Separate selection from mutation. The selection-only change is:

$\Delta p_{sel} = p_j(1 - p_j)\, \bar{s}_j(p_j)$, evaluated at the cohort-t frequency.

Balancing selection requires an internal p_j^* at which $\bar{s}_j(p_j^*) = 0$ and the selection-only recurrence is locally stable. With reversible mutation, $\Delta p_{full,j} = \Delta p_{sel,j} + \mu_j(1 - p_j) - v_j\, p_j$. If $\bar{s}_j$ is negative at every frequency, mutation may create an internal equilibrium, but persistence is mutation-selection balance rather than balancing selection. The carrier-context coefficient is:

$$\bar{s}_{j,t} = \Sigma_c \Pr(c \mid A_{j,t})\, [W(c, A_j, t) - W(c, a_j, t)] / \bar{W}_t.$$

The context index c includes both sexes, attraction profiles, expression states, path backgrounds, and fertility regimes. Female same-sex behavior and orientation are part of the allele-level portfolio; female fecundity is not the only female outcome.

Figure S2 makes the selection-only criterion explicit and shows why mutation-generated equilibrium is not balancing selection.

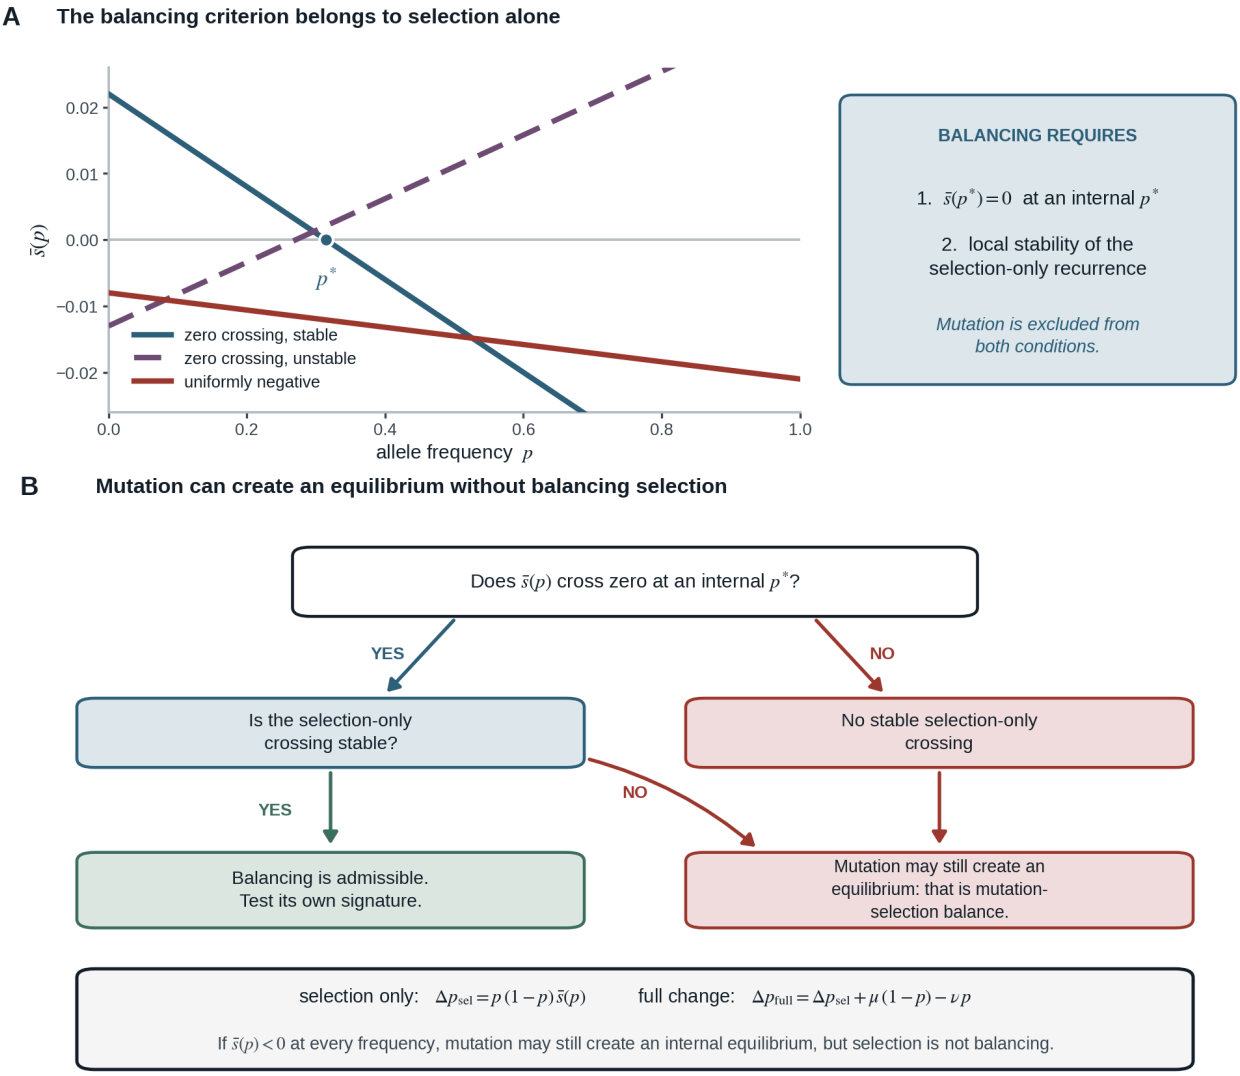


Figure S2 *When Persistence Is Balancing Selection and When It Is Not*. Panel A places the balancing criterion on the average selection coefficient alone. Panel B separates a stable internal selection-only zero from a full equilibrium created by mutation. A uniformly negative selection coefficient can coexist with an internal mutation-selection equilibrium, but it cannot support a balancing-selection claim.

S13.5 Coverage, Compensation, and the Covered-Portfolio Verdict

Let the full common-variant signal be divided into mutually exclusive, preregistered genome partitions. A bin may contain thousands of correlated effects; no small set of orthogonal components is required. A partition is selection-valid only after out-of-sample heritability or score replication, within-family validation, measurement invariance, ancestry-specific linkage-disequilibrium calibration, and fertility-regime transport. Let C be the fraction of out-of-sample SNP heritability covered by selection-valid partitions, K_C the covered expression cost, and G_C the carrier-weighted gain. The coverage and compensation criteria are those in Table 1 of the main text.

The verdict is confined to the selection-valid covered portfolio. No upper selection bound from the covered partitions is projected onto the uncovered fraction 1 − C. Reject compensation for the covered majority if its one-sided aggregate upper selection bound is below zero and a fitted mutation-selection, drift, or linkage model accounts for the observed frequencies without compensation. If the aggregate bound is below zero and no such model accounts for the observed frequencies, report the covered ledger as unbalanced; the discrepancy is not evidence of unmeasured compensation. The uncovered partitions remain unclassified. A sensitivity analysis may show what conclusions would follow under stated hypothetical residual effects, but no such assumption is labeled a bound or allowed to reverse the covered-portfolio result.

S13.6 Recurrent Developmental Incidence and Constraint

For a pregnancy-history-dependent immune mechanism, let K_t be the cohort distribution of prior male gestations and ρ_k the probability that the specified mechanism makes a causal contribution after k prior male gestations. Its marginal incidence is $q_{\mathrm{I},t} = \Sigma_k \Pr(K_t = \mathrm{k})\, \rho_k$, but the contribution to Y remains embedded in the full joint model when immune exposure interacts with inherited or regulatory variables.

Let R denote the ordinary developmental system and R' a modification that blocks the specified vulnerability. The constraint claim requires the expected fitness cost of R' to exceed the expected gain from preventing the causal outcome. A generic defense of immunity or development is insufficient. The targeted node, ordinary function, effect magnitude, and fitness consequence must be identified.

Why Critical Study 1 Requires End-to-End Simulation

A conventional single-mediation power calculation is reported here only to show why it cannot fix the design. With n = 5,000 per stratum before 15% missingness, source-to-mechanism effect .20, mechanism-to-attraction effect .08, mediator reliability .70, and two-sided alpha = .01, 1,880 of 2,000 replicates detected the indirect effect: power .940, exact 95% CI [.929, .950], with mean observed indirect effect .0111 and mean complete-case n approximately 4,250. The true indirect effect in that calculation was .016, which is below the materiality threshold in Table 1 of the main text and therefore describes a path the theory would not count. More importantly, the calculation powers one superiority test, whereas the study must sustain a weighted equivalence decision across several margins in a hold-out. Power for the former says nothing about the operating characteristics of the latter, which is why the sample size is selected by simulation of the whole decision.

The final sample size is therefore selected by end-to-end recovery simulation. Evaluate a preregistered grid around the planning design of 10,000 analyzed men per genetic-similarity stratum, adjusting in either direction as necessary. Each replicate must implement probability oversampling, inverse-probability or calibration weights, missingness, nonnormality,

measurement error, interaction, ascertainment, train-test splitting, every stratum-specific decision, all causal and predictive margins, and the mediator-outcome confounding scenarios. Simulate both portfolio and route-free data-generating processes. Adopt the smallest design with at least .90 probability of a correct retention under a portfolio generator, at least .80 probability of a correct rejection under a route-free generator, and familywise false-decision probability no greater than .05. Under the planning prevalence of .025, that design is 10,000 analyzed men per stratum, and the corresponding recruitment burden runs to between roughly 1.3 and 4.0 million invitations depending on the size of the oversample and on screening and enrollment rates.

S13.7 Critical Study 1: Do Mechanism-Distinct Causal Paths Exist?

Claims at risk. Empirical Proposition P1, supported by Definitions D3 and D4: causal path heterogeneity, level separation, and the joint-distribution architecture.

Sampling. Use probability-based screening in four genetic-similarity strata and retain inclusion probabilities for all screened and enrolled participants. Begin the simulation grid at the planning design of 10,000 analyzed men per stratum, oversample the region R for precision, and adjust in either direction until the end-to-end operating criteria above are met. Carry design weights into structural, predictive, and prevalence analyses. Include a parallel female sample for cross-sex measurement and genetic tests. The case oversample increases precision; weighted analyses recover the target population distribution only to the extent that the sampling and response model is adequate.

Measures. Assess A_{M} and A_{F} independently across repeated current and lifetime windows; behavior, fantasy, identity, partnership, and general mating effort; genome-wide genotype and sequence data; registry-verified sibships and maternal reproductive history; direct or archived pregnancy endocrine, immune, placental, and regulatory measures; family recurrence; and prospective developmental trajectories. Current maternal antibody titers are exploratory unless linked to the index pregnancy.

Model comparison. Define the adverse comparator as a route-free SCM with exactly one or two latent continuous liabilities, no source-specific mediator map, and no identified edge effects. Fit three- to five-liability route-free models as mandatory sensitivity analyses. Compare them with broad polygenicity plus measurement, independent genetic and gestational mechanisms, and the causal-portfolio graph. State each model's temporal order, mediator restrictions, conditional independences, negative controls, edge-identification claims, and external predictions. Latent classes, mixtures, and admixture models are secondary descriptions. A candidate path must pass the materiality, lower-bound, robustness, and independent-identification gates in Table 1 of the main text. Conduct full recovery simulations under skew, nonnormality, interaction, missingness, measurement error, weighting, and ascertainment.

Adverse decision rule. Use two one-sided equivalence tests with 90% confidence intervals (Lakens et al., 2018). All margins are those in Table 1 of the main text. The primary decision is made in the pooled, probability-weighted hold-out: reject the core only if fewer than two paths pass the materiality-and-identification gate and equivalence holds on the primary causal margin and on at least three of the four secondary predictive margins. The prevalence margin is not fixed independently; it is the image of the causal margin under the measurement rule and is recomputed from the fitted joint distribution. Stratum-specific results are reported as heterogeneity diagnostics rather than as gates. If stratum-specific point estimates for the primary contrast differ by more than the causal margin, the pooled verdict is withheld and the result is reported as heterogeneous. Effects from .03 to .10 are indeterminate rather than favorable to either model. A class solution without mechanism-specific validation is insufficient; a continuous outcome distribution is not adverse if two qualifying paths remain.

S13.8 Critical Study 2: Are Attraction and General Behavioral Propensity Genetically Distinguishable?

Claim at risk. Empirical Proposition P2, the genetic-layering claim, not Definition D2, which holds that attraction, behavior, identity, and partnership are noninterchangeable phenotypes.

Design. Conduct a biobank-scale genome-wide study in both sexes with multi-item repeated A_{M} and A_{F}, fantasy, identity, behavior, proportion of same-sex partners, general mating effort, risk-taking, sociosexuality, openness, pair bonding, caregiving, partner number, and completed fertility. Determine sample size by simulation from reliability, prevalence, and expected genetic covariance.

Analysis. Compare one-factor and layered genomic structural models, explicitly model indicator reliability, test measurement invariance and cross-sex covariance, and require held-out prediction of partner choice and behavior. A valuation-specific component must predict A_{M} or A_{F} after general behavioral propensity is controlled.

Adverse decision rule. Reject the genetic-layer subclaim if a one-factor genetic model is decisively preferred, attraction and behavior are predictively equivalent after reliability correction, and no attraction-specific heritable component is recoverable in either sex within a prespecified equivalence margin. The phenotypic measurement distinction remains intact.

S13.9 Critical Study 3: Is Prenatal Organization the Principal Causal Stage?

Claim at risk. Empirical Proposition P3: major biological effects on enduring attraction are established through prenatal or early perinatal organization rather than created by adult hormones or dominant postnatal social causation.

Design. Establish a prospective pregnancy cohort with fixed gestational windows for maternal and fetal steroid panels, receptor genotype and sensitivity, placental transport and enzyme

expression, immune activity, fetal growth, neonatal indices, and repeated developmental measures through adult A_{M} and A_{F}. Include postnatal social variables measured with comparable reliability and temporal resolution.

Analysis. Fit nonlinear, interaction, and path-specific models rather than a single serum-testosterone coefficient. Compare mediation of inherited and maternal sources through prenatal mechanisms against models in which postnatal variables dominate. The endocrine score must combine concentration, receptor sensitivity, placental transfer, local metabolism, and tissue-relevant proxies.

Adverse decision rule. Reject the prenatal-stage proposition if every adequately measured prenatal mechanism has an effect below the prespecified consequential bound in every powered path stratum while equally reliable postnatal variables produce robust, temporally ordered, cross-validated effects on A_{M} and A_{F}. An average null with untested nonlinear or opposing path effects is not decisive.

S13.10Critical Study 4: Does a Pregnancy-History-Dependent Maternal Immune Path Exist?

Claim at risk. The named gestational immune path, with separate decisions for NLGN4Y and the broader antigen and cellular panel.

Design. Enroll women before or during first pregnancy and collect serum, cellular, placental, and fetal-exposure measures at fixed windows in every subsequent pregnancy. Assay corrected NLGN4Y isoforms, PCDH11Y, a preregistered broad Y-antigen panel, general autoantibodies, cellular immunity, HLA type, fetal sex, pregnancy loss, microchimerism, placental transfer, and fetal growth. Follow sons to repeated A_{M} and A_{F}.

Nested predictions. The specified response must increase within women after prior male more than prior female gestations; index-pregnancy fetal exposure must prospectively predict attraction; and the immune measure must mediate a material part of the biological-older-brother association. Antigen discovery and confirmation samples must be separated.

Adverse decision rule. Failure of corrected NLGN4Y predictions rejects that antigen path. Failure across the preregistered broad antigen and cellular panel, with precise bounds on serial immunization, fetal exposure, prospective attraction, and mediation, deletes the maternal immune path. A replicated birth-order coefficient without immunization does not identify a different mechanism by default.

S13.11 Critical Study 5: Was the Covered Inherited Portfolio Nonnegative Under Natural Fertility?

Claim at risk. Empirical Proposition P4: the compensatory-maintenance arm for inherited liability.

Design and gating criteria. Use the mutually exclusive full-genome partitions defined above only after C ≥ .60 with lower bound ≥ .50 and after out-of-sample heritability, within-family validation, measurement invariance, ancestry-specific linkage-disequilibrium calibration, and fertility-regime portability are demonstrated. Primary analyses use directly genotyped natural-fertility cohorts and validated pedigrees. Historical imputation is sensitivity analysis. No existing resource meets the full gate. The feasibility program must simulate discovery samples of 250,000, 500,000, and 1,000,000 or more repeated-attraction phenotypes and the available natural-fertility validation designs, then select the lowest design that yields the required coverage and aggregate-selection precision.

Analysis. Model both sexes, expressed and nonexpressed carriers, completed fertility, offspring survival, descendant number, partnership structure, assortative mating, cohort, and gene-by-fertility-regime interaction. Estimate partition-specific selection bounds, K_C, G_C, G_C/K_C, aggregate covered selection, and allele-frequency predictions. Compare antagonistic pleiotropy, sexual antagonism, overdominance, mutation-selection balance, drift, and linkage. Report the uncovered fraction and its reasons for failing validation, but do not impute its selection distribution from the covered set.

Adverse decision rule. Reject compensatory maintenance for the covered majority if coverage and transport criteria are met, the one-sided upper bound on aggregate selection across the covered partitions is below zero, and a fitted noncompensatory process matches their observed frequencies. A negative aggregate bound that no fitted process matches is reported as an unbalanced ledger rather than as compensation. The uncovered partitions remain unclassified. A nonportable score, an assumed residual effect, or a modern partner-count association cannot rescue or invalidate the covered-portfolio verdict.

S13.12 Critical Study 6: Does a Named Placental or Regulatory Mechanism Carry a Causal Path?

Claim at risk. A specified placental or regulatory implementation. Regulatory state is not a residual class.

Design. Combine prospective placental, cord, neonatal, and endocrine measurements with genetically informed human designs and tissue-relevant cellular or animal perturbation. Adult blood and buccal samples can identify adult correlates but cannot reconstruct prenatal neural regulation. Induced pluripotent models must account for epigenetic resetting and residual source-cell memory.

Analysis. Require temporally prior prediction, mediation of a named source effect, replication in an independent prospective sample, and targeted perturbation of the developmental endpoint. Report an edge-specific effect only when its identification conditions hold; otherwise use a joint intervention, interventional analogue, or bound. Quantify any cost to ordinary differentiation, immune function, placental function, viability, or fertility.

Adverse decision rule. Delete the named mechanism if prospective association, mediation, replication, and functional perturbation all fail within prespecified bounds. An adult-tissue classifier null or a null after one reprogramming protocol rejects only that assay.

S13.13Two Subsidiary Tests

Comparative sex-discrimination claim. Across species, standardized rates and contexts of same-sex behavior should be modeled against encounter rate, sex ratio, mating cost, opportunity cost, acceptance thresholds, sexual-signal overlap, and vomeronasal-pathway integrity while controlling phylogeny and observation effort (Lerch, 2026). The TRPC2 claim fails if pathway integrity and developmentally timed perturbation add no prediction beyond ecology and adult sensory lesions.

Kin-investment claim. In several kin-dense societies, follow androphilic men, childless gynephilic men, and childless women; record actual time and material transfers; and estimate effects on recipient survival and completed reproduction. The channel is evolutionarily negligible if transfers do not generate an inclusive-fitness return approaching forgone direct reproduction.

S13.14Graded Inferential Consequences

The consequences are asymmetric. Fewer than two paths passing the materiality-and-identification gate, together with a route-free one- or two-liability SCM that is equivalent on the primary causal margin and a majority of the secondary margins in the pooled hold-out, rejects the causal-portfolio core. Genetic equivalence of attraction and behavioral propensity removes genetic layering but not measurement distinctness. A negative selection-valid covered portfolio rejects compensation for that majority while leaving uncovered partitions unclassified and causal heterogeneity potentially intact. Failure of the broad maternal immune path deletes that path. Failure of a named regulatory mechanism deletes that mechanism. One peripheral assay null remains assay-specific.

No study has implemented the complete designs and decision rules above. Components of every design exist, but Critical Study 5 in particular requires data not presently assembled. The theory is decisive because it states which complete patterns overturn the architecture, which force reconstruction, which delete a path, which are long-horizon rather than currently runnable, and which merely invalidate a measurement implementation.

Supplementary References

Abé, C., Lebedev, A., Zhang, R., Jonsson, L., Bergen, S. E., Ingvar, M., Landén, M., & Rahman, Q. (2021). Cross-sex shifts in two brain imaging phenotypes and their relation to polygenic scores for same-sex sexual behavior: A study of 18,645 individuals from the UK Biobank. *Human Brain Mapping*, *42*(7), 2292–2304. https://doi.org/10.1002/hbm.25370

Ablaza, C., Kabátek, J., & Perales, F. (2022). Are sibship characteristics predictive of same sex marriage? An examination of fraternal birth order and female fecundity effects in population-level administrative data from the Netherlands. *The Journal of Sex Research*, *59*(6), 671–683. https://doi.org/10.1080/00224499.2021.1974330

Alanko, K., Santtila, P., Harlaar, N., Witting, K., Varjonen, M., Jern, P., Johansson, A., von der Pahlen, B., & Sandnabba, N. K. (2010). Common genetic effects of gender atypical behavior in childhood and sexual orientation in adulthood: A study of Finnish twins. *Archives of Sexual Behavior*, *39*(1), 81–92. https://doi.org/10.1007/s10508-008-9457-3

Allen, L. S., & Gorski, R. A. (1992). Sexual orientation and the size of the anterior commissure in the human brain. *Proceedings of the National Academy of Sciences of the United States of America*, *89*(15), 7199–7202. https://doi.org/10.1073/pnas.89.15.7199

Avin, C., Shpitser, I., & Pearl, J. (2005). Identifiability of path-specific effects. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence* (pp. 357–363).

Bailey, J. M., & Pillard, R. C. (1991). A genetic study of male sexual orientation. *Archives of General Psychiatry*, *48*(12), 1089–1096.

Bailey, J. M., & Zucker, K. J. (1995). Childhood sex-typed behavior and sexual orientation: A conceptual analysis and quantitative review. *Developmental Psychology*, *31*(1), 43–55. https://doi.org/10.1037/0012-1649.31.1.43

Bailey, J. M., Dunne, M. P., & Martin, N. G. (2000). Genetic and environmental influences on sexual orientation and its correlates in an Australian twin sample. *Journal of Personality and Social Psychology*, *78*(3), 524–536. https://doi.org/10.1037/0022-3514.78.3.524

Bailey, J. M., Pillard, R. C., Dawood, K., Miller, M. B., Farrer, L. A., Trivedi, S., & Murphy, R. L. (1999). A family history study of male sexual orientation using three independent samples. *Behavior Genetics*, *29*(2), 79–86.

Bailey, J. M., Vasey, P. L., Diamond, L. M., Breedlove, S. M., Vilain, E., & Epprecht, M. (2016). Sexual orientation, controversy, and science. *Psychological Science in the Public Interest*, *17*(2), 45–101. https://doi.org/10.1177/1529100616637616

Bailey, N. W., & Zuk, M. (2009). Same-sex sexual behavior and evolution. *Trends in Ecology & Evolution*, *24*(8), 439–446. https://doi.org/10.1016/j.tree.2009.03.014

Barthes, J., Godelle, B., & Raymond, M. (2013). Human social stratification and hypergyny: Toward an understanding of male homosexual preference. *Evolution and Human Behavior*, *34*(3), 155–163. https://doi.org/10.1016/j.evolhumbehav.2013.01.001

Bauer, D. J., & Curran, P. J. (2003). Distributional assumptions of growth mixture models: Implications for overextraction of latent trajectory classes. *Psychological Methods*, *8*(3), 338–363. https://doi.org/10.1037/1082-989X.8.3.338

Blanchard, R., & Bogaert, A. F. (2004). Proportion of homosexual men who owe their sexual orientation to fraternal birth order: An estimate based on two national probability samples. *American Journal of Human Biology*, *16*(2), 151–157. https://doi.org/10.1002/ajhb.20006

Blanchard, R., Kabátek, J., & Zdaniuk, B. (2025). A current approach to logistic regression analysis of birth order and sexual orientation. *Archives of Sexual Behavior*, *54*, 4013–4022. https://doi.org/10.1007/s10508-025-03275-3

Blanchard, R., Krupp, J., VanderLaan, D. P., Vasey, P. L., & Zucker, K. J. (2020). A method yielding comparable estimates of the fraternal birth order and female fecundity effects in male homosexuality. *Proceedings of the Royal Society B: Biological Sciences*, *287*(1923), 20192907. https://doi.org/10.1098/rspb.2019.2907

Blumenthal, S. A., & Young, L. J. (2023). The neurobiology of love and pair bonding from human and animal perspectives. *Biology*, *12*(6), 844.

Bobrow, D., & Bailey, J. M. (2001). Is male homosexuality maintained via kin selection? *Evolution and Human Behavior*, *22*, 361–368. https://doi.org/10.1016/S1090-5138(01)00074-5

Bocklandt, S., Horvath, S., Vilain, E., & Hamer, D. H. (2006). Extreme skewing of X chromosome inactivation in mothers of homosexual men. *Human Genetics*, *118*(6), 691–694. https://doi.org/10.1007/s00439-005-0119-4

Bogaert, A. F., Skorska, M. N., Wang, C., Gabrie, J., MacNeil, A. J., Hoffarth, M. R., VanderLaan, D. P., Zucker, K. J., & Blanchard, R. (2018). Male homosexuality and maternal immune responsivity to the Y-linked protein NLGN4Y. *Proceedings of the National Academy of Sciences*, *115*(2), 302–306. https://doi.org/10.1073/pnas.1705895114

Bogaert, A. F., Skorska, M. N., Wang, C., Gabrie, J., MacNeil, A. J., Hoffarth, M. R., VanderLaan, D. P., Zucker, K. J., & Blanchard, R. (2022). Correction to Supporting Information for Bogaert et al., "Male homosexuality and maternal immune responsivity to the Y-linked protein NLGN4Y." *Proceedings of the National Academy of Sciences*, *119*(34), e2212767119. https://doi.org/10.1073/pnas.2212767119

Byne, W., Tobet, S., Mattiace, L. A., Lasco, M. S., Kemether, E., Edgar, M. A., Morgello, S., Buchsbaum, M. S., & Jones, L. B. (2001). The interstitial nuclei of the human anterior hypothalamus: An investigation of variation with sex, sexual orientation, and HIV status. *Hormones and Behavior*, *40*(1), 86–92. https://doi.org/10.1006/hbeh.2001.1680

Camperio Ciani, A., & Pellizzari, E. (2012). Fecundity of paternal and maternal non-parental female relatives of homosexual and heterosexual men. *PLoS ONE*, *7*(12), e51088. https://doi.org/10.1371/journal.pone.0051088

Camperio Ciani, A., Cermelli, P., & Zanzotto, G. (2008). Sexually antagonistic selection in human male homosexuality. *PLoS ONE*, *3*(6), e2282. https://doi.org/10.1371/journal.pone.0002282

Cantor, J. M., Blanchard, R., Paterson, A. D., & Bogaert, A. F. (2002). How many gay men owe their sexual orientation to fraternal birth order? *Archives of Sexual Behavior*, *31*(1), 63–71. https://doi.org/10.1023/A:1014031201935

Cinelli, C., & Hazlett, C. (2020). Making sense of sensitivity: Extending omitted variable bias. *Journal of the Royal Statistical Society: Series B*, *82*, 39–67. https://doi.org/10.1111/rssb.12348

Clive, J., Flintham, E., & Savolainen, V. (2023). Same-sex sociosexual behaviour is widespread and heritable in male rhesus macaques. *Nature Ecology & Evolution*, *7*, 1287–1301. https://doi.org/10.1038/s41559-023-

02111-y

Clive, J., Wisden, W., & Savolainen, V. (2021). The de-scent of sexuality: Should we smell a rat? *Archives of Sexual Behavior*, *50*, 2283–2288. https://doi.org/10.1007/s10508-019-01591-z

Felesina, T. (2026). Understanding the evolutionary history of human male and female same-sex sexual behavior requires a greater appreciation of its shared genetic architecture. *Archives of Sexual Behavior*, *55*(4), 1451–1453. https://doi.org/10.1007/s10508-026-03478-2

Felesina, T., & Zietsch, B. P. (2024). The desirable dad hypothesis: Male same-sex attraction as the product of selection for paternal care via antagonistic pleiotropy. *Archives of Sexual Behavior*, *53*(5), 1731–1745.

Felesina, T., & Zietsch, B. P. (2025). Emerging insights into the genetics and evolution of human same-sex sexual behavior. *Trends in Genetics*, *41*(5), 402–411. https://doi.org/10.1016/j.tig.2024.12.005

Folkierska-Żukowska, M., & Dragan, W. Ł. (2025). Biodevelopmental correlates of sexual orientation in men: Evidence from a Polish sample. *Archives of Sexual Behavior*, *54*(8), 2825–2838. https://doi.org/10.1007/s10508-024-03018-w

Fořt, J., Flegr, J., Kuba, R., & Kaňková, Š. (2024). Fertility of Czech gay and straight men, women, and their relatives: Testing the sexually antagonistic gene hypothesis. *Archives of Sexual Behavior*, *53*(5), 1747–1761.

Frigerio, A., Ballerini, L., & Valdés Hernández, M. (2021). Structural, functional, and metabolic brain differences as a function of gender identity or sexual orientation: A systematic review of the human neuroimaging literature. *Archives of Sexual Behavior*, *50*, 3329–3352. https://doi.org/10.1007/s10508-021-02005-9

Ganna, A., Verweij, K. J. H., Nivard, M. G., Maier, R., Wedow, R., Busch, A. S., Abdellaoui, A., Guo, S., Sathirapongsasuti, J. F., 23andMe Research Team, Lichtenstein, P., Lundström, S., Långström, N., Auton, A., Harris, K. M., Beecham, G. W., Martin, E. R., Sanders, A. R., Perry, J. R. B., … Zietsch, B. P. (2019). Large-scale GWAS reveals insights into the genetic architecture of same-sex sexual behavior. *Science*, *365*(6456), eaat7693. https://doi.org/10.1126/science.aat7693

Gavrilets, S., & Rice, W. R. (2006). Genetic models of homosexuality: Generating testable predictions. *Proceedings of the Royal Society B: Biological Sciences*, *273*, 3031–3038. https://doi.org/10.1098/rspb.2006.3684

Hamer, D. H. (1999). Genetics and male sexual orientation [Technical comment with response by G. Rice, N. Risch, & G. Ebers]. *Science*, *285*(5429), 803a. https://doi.org/10.1126/science.285.5429.803a

Hamer, D. H., Hu, S., Magnuson, V. L., Hu, N., & Pattatucci, A. M. L. (1993). A linkage between DNA markers on the X chromosome and male sexual orientation. *Science*, *261*(5119), 321–327.

Hernán, M. A., & Robins, J. M. (2020). *Causal inference: What if*. Chapman & Hall/CRC.

Hu, S., Pattatucci, A. M. L., Patterson, C., Li, L., Fulker, D. W., Cherny, S. S., Kruglyak, L., & Hamer, D. H. (1995). Linkage between sexual orientation and chromosome Xq28 in males but not in females. *Nature Genetics*, *11*(3), 248–256. https://doi.org/10.1038/ng1195-248

Hu, S.-H., Li, H.-M., Yu, H., Liu, Y., Liu, C.-X., Zuo, X.-B., Lu, J., Jiang, J.-J., Xi, C.-X., Huang, B.-C., Xu, H.-J., Hu, J.-B., Lai, J.-B., Huang, M.-L., Liu, J.-N., Xu, D.-G., Guo, X.-C., Wu, W., Wu, X., … Xu, Y. (2021). Discovery of new genetic loci for male sexual orientation in Han population. *Cell Discovery*, *7*, 103. https://doi.org/10.1038/s41421-021-00341-7

Imai, K., Keele, L., & Tingley, D. (2010). A general approach to causal mediation analysis. *Psychological Methods*, *15*(4), 309–334. https://doi.org/10.1037/a0020761

Kendler, K. S., Thornton, L. M., Gilman, S. E., & Kessler, R. C. (2000). Sexual orientation in a U.S. national sample of twin and nontwin sibling pairs. *American Journal of Psychiatry*, *157*(11), 1843–1846. https://doi.org/10.1176/appi.ajp.157.11.1843

Kirkpatrick, R. C. (2000). The evolution of human homosexual behavior. *Current Anthropology*, *41*(3), 385–413.

Lakens, D., Scheel, A. M., & Isager, P. M. (2018). Equivalence testing for psychological research: A tutorial. *Advances in Methods and Practices in Psychological Science*, *1*, 259–269. https://doi.org/10.1177/2515245918770963

Långström, N., Rahman, Q., Carlström, E., & Lichtenstein, P. (2010). Genetic and environmental effects on same-sex sexual behavior: A population study of twins in Sweden. *Archives of Sexual Behavior*, *39*(1), 75–80. https://doi.org/10.1007/s10508-008-9386-1

Lerch, B. A. (2026). The evolution of same-sex sexual behavior: Using old theory to answer new questions. *Integrative and Comparative Biology*, *66*, Article icag015. https://doi.org/10.1093/icb/icag015

LeVay, S. (1991). A difference in hypothalamic structure between heterosexual and homosexual men. *Science*, *253*(5023), 1034–1037.

Miller, E. M. (2000). Homosexuality, birth order, and evolution: Toward an equilibrium reproductive economics of homosexuality. *Archives of Sexual Behavior*, *29*(1), 1–34.

Mustanski, B. S., DuPree, M. G., Nievergelt, C. M., Bocklandt, S., Schork, N. J., & Hamer, D. H. (2005). A genomewide scan of male sexual orientation. *Human Genetics*, *116*(4), 272–278. https://doi.org/10.1007/s00439-004-1241-4

Ngun, T. C., Guo, W., Ghahramani, N. M., Purkayastha, K., Conn, D., Sanchez, F. J., Bocklandt, S., Zhang, M., Ramirez, C. M., Pellegrini, M., & Vilain, E. (2015, October 8). *A novel predictive model of sexual orientation using epigenetic markers* (Abstract No. 95) [Conference presentation]. American Society of Human Genetics 2015 Annual Meeting, Baltimore, MD, United States.

Pearl, J. (2009). *Causality: Models, reasoning, and inference* (2nd ed.). Cambridge University Press.

Pfau, D. R., Jordan, C. L., & Breedlove, S. M. (2021a). The de-scent of sexuality: Did loss of a pheromone signaling protein permit the evolution of same-sex sexual behavior in primates? *Archives of Sexual Behavior*, *50*(6), 2267–2276. https://doi.org/10.1007/s10508-018-1377-2

Pfau, D. R., Jordan, C. L., & Breedlove, S. M. (2021b). Response to commentaries: Sniffing out consensus on the evolution of primate same-sex sexual behavior. *Archives of Sexual Behavior*, *50*(6), 2317–2320. https://doi.org/10.1007/s10508-021-02078-6

Rahman, Q., & Wilson, G. D. (2003). Born gay? The psychobiology of human sexual orientation. *Personality and Individual Differences*, *34*(8), 1337–1382. https://doi.org/10.1016/S0191-8869(02)00140-X

Rice, G., Anderson, C., Risch, N., & Ebers, G. (1999). Male homosexuality: Absence of linkage to microsatellite markers at Xq28. *Science*, *284*(5414), 665–667. https://doi.org/10.1126/science.284.5414.665

Robins, J. M., & Greenland, S. (1992). Identifiability and exchangeability for direct and indirect effects. *Epidemiology*, *3*(2), 143–155. https://doi.org/10.1097/00001648-199203000-00013

Roselli, C. E., Reddy, R. C., & Kaufman, K. R. (2011). The development of male-oriented behavior in rams. *Frontiers in Neuroendocrinology*, *32*(2), 164–169. https://doi.org/10.1016/j.yfrne.2010.12.007

Sanders, A. R., Martin, E. R., Beecham, G. W., Guo, S., Dawood, K., Rieger, G., Badner, J. A., Gershon, E. S., Krishnappa, R. S., Kolundzija, A. B., Duan, J., Gejman, P. V., & Bailey, J. M. (2015). Genome-wide scan demonstrates significant linkage for male sexual orientation. *Psychological Medicine*, *45*(7), 1379–1388. https://doi.org/10.1017/S0033291714002451

Savic, I., & Lindström, P. (2008). PET and MRI show differences in cerebral asymmetry and functional connectivity between homo- and heterosexual subjects. *Proceedings of the National Academy of Sciences*, *105*, 9403–9408.

Savic, I., Berglund, H., & Lindström, P. (2005). Brain response to putative pheromones in homosexual men. *Proceedings of the National Academy of Sciences*, *102*, 7356–7361.

Semenyna, S. W., Gómez Jiménez, F. R., VanderLaan, D. P., & Vasey, P. L. (2023). Male androphilia, fraternal birth order, and female fecundity in Samoa: A 10-y retrospective. *Proceedings of the National Academy of Sciences*, *120*(50), e2313284120. https://doi.org/10.1073/pnas.2313284120

Shapley, L. S. (1953). A value for n-person games. In H. W. Kuhn & A. W. Tucker (Eds.), *Contributions to the theory of games* (Vol. 2, pp. 307–317). Princeton University Press.

Song, S., & Zhang, J. (2024). Genetic variants underlying human bisexual behavior are reproductively advantageous. *Science Advances*, *10*(1), eadj6958. https://doi.org/10.1126/sciadv.adj6958

Swaab, D. F., & Hofman, M. A. (1990). An enlarged suprachiasmatic nucleus in homosexual men. *Brain Research*, *537*(1–2), 141–148. https://doi.org/10.1016/0006-8993(90)90350-K

Swift-Gallant, A., Aung, T., Rosenfield, K., Dawood, K., & Puts, D. (2023). Organizational effects of gonadal hormones on human sexual orientation. *Adaptive Human Behavior and Physiology*, *9*, 344–370. https://doi.org/10.1007/s40750-023-00226-x

Swift-Gallant, A., Coome, L. A., Aitken, M., Monks, D. A., & VanderLaan, D. P. (2019). Evidence for distinct biodevelopmental influences on male sexual orientation. *Proceedings of the National Academy of Sciences*, *116*(26), 12787–12792. https://doi.org/10.1073/pnas.1809920116

VanderLaan, D. P., Skorska, M. N., Peragine, D. E., & Coome, L. A. (2022). Carving the biodevelopment of same-sex sexual orientation at its joints. In D. P. VanderLaan & W. I. Wong (Eds.), *Gender and sexuality development: Contemporary theory and research* (pp. 491–538). Springer. https://doi.org/10.1007/978-3-030-84273-4_17

VanderWeele, T. J. (2015). *Explanation in causal inference: Methods for mediation and interaction*. Oxford University Press.

VanderWeele, T. J., & Ding, P. (2017). Sensitivity analysis in observational research: Introducing the E-value. *Annals of Internal Medicine*, *167*, 268–274. https://doi.org/10.7326/M16-2607

VanderWeele, T. J., Vansteelandt, S., & Robins, J. M. (2014). Effect decomposition in the presence of an exposure-induced mediator-outcome confounder. *Epidemiology*, *25*(2), 300–306. https://doi.org/10.1097/EDE.0000000000000034

Vasey, P. L., & VanderLaan, D. P. (2010). Avuncular tendencies and the evolution of male androphilia in Samoan fa'afafine. *Archives of Sexual Behavior*, *39*, 821–830. https://doi.org/10.1007/s10508-008-9404-3

Vilsmeier, J. K., Kossmeier, M., Voracek, M., & Tran, U. S. (2023). The fraternal birth-order effect as a statistical artefact: Convergent evidence from probability calculus, simulated data, and multiverse meta-analysis. *PeerJ*, *11*, e15623.

Zietsch, B. P., Morley, K. I., Shekar, S. N., Verweij, K. J. H., Keller, M. C., Macgregor, S., Wright, M. J., Bailey, J. M., & Martin, N. G. (2008). Genetic factors predisposing to homosexuality may increase mating success in heterosexuals. *Evolution and Human Behavior*, *29*(6), 424–433.

Zietsch, B. P., Sidari, M. J., Abdellaoui, A., Maier, R., Långström, N., Guo, S., Beecham, G. W., Martin, E. R., Sanders, A. R., & Verweij, K. J. H. (2021). Genomic evidence consistent with antagonistic pleiotropy may help explain the evolutionary maintenance of same-sex sexual behaviour in humans. *Nature Human Behaviour*, *5*(9), 1251–1258.

Recommended Citation

Moore, T. M. (2026). *The evolutionary persistence of male androphilia: A causal-portfolio theory of developmental paths and selection accounting*. Manuscript retrieved from https://mooremetrics.com/homosexuality. https://doi.org/10.5281/zenodo.23789873