Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

When Time Refutes: Minimum Generation Length and Claims of Joint Hashemite Descent; Exclusion tests on the claimant clades under J-FGC11, from Hashim to 'Adnan

Sultan Faleh Alshaabi

alshaabi@hotmail.com

Independent Researcher

Research Article

Keywords: Y chromosome, genetic genealogy, generation length, coalescence time, exclusion test, haplogroup J1

Posted Date: November 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11727955/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

When Time Refutes: Minimum Generation Length and Claims of Joint Hashemite Descent

Exclusion tests on the claimant clades under J-FGC11, from Hashim to ‘Adnan

Sultan Faleh Alshaabi

Independent Researcher, Medina, Saudi Arabia

Corresponding author: alshaabi@hotmail.com

ORCID: 0009-0001-4696-1662

Abstract

Descent claims in genetic genealogy usually rest on membership in named clusters, inheriting whatever identity the cluster was assigned. This paper proposes a chronological test that identifies no cluster with any ancestor: if two lineages descend patrilineally from one ancestor, their most recent common ancestor cannot predate him. Given the meeting age M, the reference-generation date T_0 and the generation count k, the data set a lower bound on mean generation length, $g \geq (M - T_0) \div k$; above a declared ceiling of 35 years, joint descent is excluded.

A declared rule draws the sides from public data: every clade under J-FGC11 no older than Hashim’s time in which a sample carries one of four descent names, yielding sixteen units, 75 records and 120 pairs. Exclusion at Hashim is guaranteed by this construction; the findings are the grades. Meeting ages are resolved under three published mutation rates; an exclusion is firm only at probability ≥ 0.95 under all three.

Five units part from the rest about 2,900–3,100 years ago and are excluded at every rung from Hashim to ‘Adnan. The remaining eleven fall in four groups that part about 2,200 years ago and cannot share an ancestor from Hashim to Ghalib. Of the 120 pairs, 115 are firmly excluded at Hashim. Five are not: the three among J-ZS4753, J-BY258 and J-FGC11743 under J-FGC4302 (probable exclusion), and one pair each inside J-FGC7944 and J-KM2 (undecided). The test establishes no descent and does not locate where sides meet.

Keywords: Y chromosome; genetic genealogy; generation length; coalescence time; exclusion test; haplogroup J1

1. The problem

A number of genetic lineages are today claimed for named ancestors: for Hashim, Quraysh, Kinana, Mudrika or those above them. These claims are current in the genetic-genealogy community and are stated publicly, so they are not inferences of this paper.

Such claims are usually contested with classificatory tools: does this lineage fall inside the cluster bearing the ancestor’s name, or outside it? To which it may be answered that the naming is a human decision, and that a cluster’s boundary shifts when a sample is added or a classification revised.

An attempt to repair this runs into a subtler difficulty. Suppose one said: let us measure every lineage against the line that carries the Prophet’s genealogy. That would be sound in itself; but the genetic tree

carries no names, so the identification of that line cannot come from the tree. It comes from the very declarations and labels that were rejected at the outset. Whoever designates a reference line has admitted through the window what he expelled through the door, and has made that line untestable by definition rather than by evidence.

This paper therefore designates no reference line. The quantity it measures is the coalescence time between the sides under study, a fact of the tree that depends on nobody's choice, and the verdict it gives is exclusion. Where a reference helps the reading, every side is taken in turn as holding the measure (Section 3.5), so that none is privileged and the verdicts can be shown not to depend on the choice.

The Prophet's genealogy here is the measure, not the measured. It is attested by mass transmission and is not subjected to these tools. It enters the computation in one way only: the time T_0 of his generation and the number k of generations above it. Which genetic clade carries it is a question the body of this paper does not answer.

2. Related work

Cross-cultural and genomic estimates place the mean paternal generation interval near 30–32 years (Fenner 2005; Wang et al. 2023). Wang et al. (2023), using changes in the mutation spectrum across whole-genome data, estimated an average male generation time of 30.7 years over the past 250,000 years. These two estimates are the source of the ceiling declared in this paper. They are entirely independent of the tree under study, so they introduce into it no assignment and no classification.

Statistical work on Y-chromosome dating has focused mainly on refining estimation of the time to the most recent common ancestor (TMRCA). McDonald (2021) developed improved probabilistic models incorporating Y-chromosome single-nucleotide polymorphism (Y-SNP) and short tandem repeat (Y-STR) data together with historical constraints, and recent analyses have examined whether plausible variation in male generation time can account for observed branch-length variation, arguing that it cannot be the sole explanation (Swiel et al. 2025).

These studies treat generation length as a parameter affecting age estimates. They do not invert the relation to derive a minimum generation length from a claimed ancestral date. That inversion is the subject of this paper.

Earlier genetic studies of historical descent claims, most notably those on the Jewish priesthood, tested shared ancestry through cluster membership and modal haplotypes (Skorecki et al. 1997; Hammer et al. 2009). A recent high-resolution next-generation sequencing (NGS) re-examination has refined the picture into multiple distinct Cohen branches while still relying on clade membership (Lipson et al. 2025).

The present paper replaces membership with a pairwise chronological lower bound on generation length, read as exclusion only and under every possible internal reference. It requires no designated reference lineage and no classificatory label.

3. The criterion

3.1 Definitions

Let T_0 denote the time of the Prophet's generation in years before present. His birth in 570 CE and a present of 2026 give $T_0 = 1456$ years. It is one of two quantities taken from the historical record, the other being the generation count k; the paper takes nothing else from it. The birth year is itself approximate to within about five years, and its effect on the computation is very slight.

Let k denote the number of generations between the claimed ancestor and the Prophet's generation along the attested line: Hashim 3, Qusayy 5, Murra 7, Ka'b 8, Lu'ayy 9, Ghalib 10, Fihr 11, al-Nadr 13, Kinana 14, Mudrika 16, Ilyas 17, Mudar 18, 'Adnan 21. The line is taken as transmitted in the standard sources, which agree on it from 'Adnan downward (Ibn Hisham, al-Sira; Ibn Sa'd, al-Tabaqat; Ibn Hazm, Jamhara).

Two points of naming should be settled. Clades are referred to throughout by their SNP identifiers alone; the names of ancestors denote rungs of the ladder, that is values of k, and are never attached to a node of the tree. And for "Quraysh" there are two well-known views as to which ancestor the name attaches: Fihr b. Malik (k = 11) and al-Nadr b. Kinana (k = 13). Every rung is named after its ancestor, so the rung k = 11 is labelled Fihr; the alternative view is examined in Section 8.2.

A **side** is one of the clades under study together with all of its subclades. Let M denote the coalescence time of two sides, the age of their most recent common ancestor read from the dated tree, and g the mean generation length. The node at which two sides meet is called their **meeting node**.

3.2 The constraint is one-sided

If two lineages descend from a single ancestor standing k generations above the Prophet's generation, their coalescence must fall at that ancestor or later, since two descendants of a man cannot part before he exists. That ancestor's time is $T_0 + k \cdot g$. From the constraint

$$M \leq T_0 + k \cdot g$$

it follows that

$$g \geq (M - T_0) \div k$$

This is a constraint from one side only. Coalescence time yields a lower bound on generation length and no upper bound. A clade's formation time must therefore not be used as an upper bound on the ancestor; that common error manufactures a closed interval where none exists. A pair coalescing more recently than T_0 produces no constraint at all.

The entailment is logical, not probabilistic: granted the premises, the conclusion follows. The uncertainty lies in the premises: in the estimated M, in the k taken from the record, and in the declared ceiling on generation length. Each verdict is therefore reported with the adjustment that would overturn it.

Coalescence time is not the founding ancestor's time. It is the time of the most recent common ancestor of those who happen to have been sampled. For the quantity used here this difference does not bite: the common ancestor of samples drawn from two distinct sides is their point of divergence itself, and sparse sampling within either side does not change which node that is. Sampling does bear on how well the node is dated and on whether a polytomy such as J-FGC23310 is resolved (Section 9), so no claim of

conservatism is made from sampling. Two things alone protect the verdicts: the one-sidedness of the constraint, and the generosity of the declared ceiling toward the claimant.

3.3 Why pairs and not individuals

First, a pairwise constraint needs no reference. The statement “these two do not share that ancestor” is true or false regardless of which lineage anyone calls the correct one, since their coalescence time is the same either way. The statement “this one does not descend from him”, taken alone, requires a line to measure against.

Second, a pairwise coalescence is a deep node supported by many mutations and many samples, unlike a clade’s own age, which depends on which of its carriers happen to be sampled. Section 7 shows how disagreement between databases varies with depth.

If two sides are said not to share an ancestor, that does not identify which of them fails; it follows only that at least one does not descend from him.

3.4 The form of the verdict

The only verdict is exclusion. The exclusion of two sides is a property of the two alone. Adding a third claim, removing one, or replacing a side’s representative by another of its subclades cannot lift it, since two distinct sides diverge at the same node whichever of their descendants are sampled. Non-exclusion lacks this property. Two claims lying close together in the tree remain compatible down to the age of their divergence, whatever their truth, and any set assembled from non-exclusions changes with the claims admitted to it. Non-exclusion is therefore reported as the silence of the instrument and given no weight.

Verdicts are counted on meeting nodes. Many pairs pass through few meeting nodes: the 120 pairs of the open claim pass through nine levels, and the six pairs of the four clades followed in detail through three nodes. A verdict on a node is one verdict however many pairs pass through it. Every count in Section 7 is made per node.

An exclusion is graded by one rule. Where the meeting age can be resolved from branch SNP counts (Section 7.5), an exclusion is called firm when the probability that the divergence predates the ancestor is at least 0.95 under each of three published Y-SNP mutation rates, probable when it is at least 0.80 under each, and undecided otherwise. The figure of 0.95 is a declared decision threshold, not the probability that a verdict is true. The node at which L859 meets the other sides lies so far from every threshold that no published rate approaches it (Section 8.5), and it is graded from the central estimates of both databases. The central estimates of the two databases, shown in several tables as E (excluded) or a dash, and the stability figures of Section 7.4 are descriptive: they show where each database places a node, and the grade alone is the verdict.

The rule can be written as a single condition. A pair is excluded at rung k when

$$M(5\%) > T_0 + k \cdot g(\max)$$

where M(5%) is the lower bound of a one-sided 95% interval for the meeting age and g(max) the declared ceiling. An exclusion is firm when the condition holds under each of the three published rates. This is the probability rule above in another form: the verdict rests on the lower bound of the age and never on its central value. The intervals are shown in Figure 2.

3.5 Multiple internal references

Each side is taken in turn as holding the measure. For the purpose of reading, it is supposed that the attested line runs through some subclade of that side, and the other sides are tested against that subclade. No side is chosen over another; all are taken. Any dated subclade of a side is an eligible reference.

The verdict is the same for every eligible reference, by a property of trees: if X and Y are disjoint clades, the most recent common ancestor of any descendant of X and any descendant of Y is the most recent common ancestor of X and Y themselves. The meeting age of a reference in side X with any subclade of side Y is therefore the meeting age of the two sides. Table A5 checks that the reconstructed trees satisfy this. Each verdict therefore reads: if the measure lies anywhere in side X, side Y is excluded at the rungs stated.

This reading does not identify which side holds the measure. That identification would require a premise from outside the tree, and the paper uses none.

3.6 How the sides are fixed

The sides are not chosen by the author. A claim is read from the public record by a fixed rule: a sample is a claimant if its displayed name or its stated earliest paternal ancestor contains one of four descent names, Hashimi, Hasani, Husayni or Idrisi, in Arabic or Latin script, including the collective forms Banu Hashim, al-Hawashim and al-Adarisa. Two public sources are searched: the sample names shown in Source A, and the member lists of nine public group projects, most of them hosted at FamilyTreeDNA. The names serve only to fix which clades are tested. They enter no computation, and no verdict depends on whether any of them is true.

A **unit** is the largest clade containing a claimant whose own coalescence in Source B is no older than the Hashim threshold (1561 years). The definition follows from the criterion: a clade claimed as a whole for Hashim must be younger than he is. A claimant whose terminal SNP lies above every such clade defines no unit and is set aside. Each unit is a side. By the property of Section 3.5 every lineage inside a unit takes the unit's verdicts, so one claimant is enough for a unit to enter, and the number of claimants in it has no bearing on any verdict.

The definition has a consequence that must be stated before any result. Two distinct units cannot lie one inside the other, so they meet at a node older than both, and since each is the largest clade below the threshold, that node is older than 1561 years in Source B. Exclusion at Hashim by the central estimates of Source B therefore holds for every pair of units by construction and is not a finding. Nor is it rare: of all pairs of samples drawn at random from the Source B tree, 96.4% meet above the Hashim threshold and 93.3% above that of 'Adnan. What the test can add is the grade, that is whether the exclusion survives the uncertainty of the meeting age, and how far up the ladder it reaches.

The rule yields sixteen units, and all sixteen are carried into the tables. The search was automatic and no record was dropped on reading, because any criterion for dropping would be the author's. Some records plainly use one of the four words otherwise than as a claim of descent: a family name, or a tribal name attached to a tribe of other stock. They are described in Section 9; keeping them adds units and pairs and removes no exclusion.

4. Data

Two dating databases were used, and coalescence times were computed on each separately. Values were not pooled and no average taken:

- Source A — TheYtree: the dated J-FGC11 tree (3,875 downstream branches; in the page used, 3,355 dated nodes and 3,732 samples, with parentage preserved). It publishes no version number and no update date; accessed November 2026.
- Source B — YFull YTree v14.07.00 (2 November 2026; YFull 2026a): the complete J-FGC11 tree, 4,559 dated nodes and 7,025 samples.

The two databases are maintained independently. TheYtree states that it keeps its own phylogenetic tree, separate from YFull's. They may nevertheless draw on overlapping sample data, so their agreement cannot be read as a bound on error: two databases may agree on one figure and be wrong together. The dispersion between them is reported only as a description of how stable a verdict is between two particular databases.

Table 1 The claimant units fixed by the rule of Section 3.6: position, own coalescence in each source, samples in Source B, and records carrying one of the four names

Unit	Position	Coalescence B	Coalescence A	Samples (B)	Records	Names in the records
J-ZS4753	under J-FGC4302	1400	1550	110	24	Idrisi 19, Hashimi 11, Hasani 7, Husayni 1
J-BY258	under J-FGC4302, within J-FGC4309	1350	1450	43	2	Idrisi 1, Husayni 1
J-FGC11743	under J-FGC4302, within J-FGC10092	1350	1490	15	1	Hasani 1
J-ZS11906	under J-FGC4316, outside J-FGC4302	550	550	6	1	Idrisi 1
J-FGC51966	under J-FGC7, outside J-FGC4316	1400	1400	10	1	Hashimi 1
J-FT195589	under J-FGC7944	800	—	7	5	Idrisi 4, Hashimi 1
J-Y207222	under J-FGC7944, beside J-FT195589	750	—	3	1	Hashimi 1
J-Y8537	under J-KM2	1400	1790 *	1,098	23	Hashimi 21, Idrisi 2
J-FT198133	under J-KM2, beside J-Y8537	1200	850	3	1	Idrisi 1
J-FTE99194	under J-KM2, outside J-FGC5	500	490	1	1	Hashimi 1
J-FT38741	under J-FGC1721, outside J-FGC23310	1250	—	2	1	Idrisi 1
J-L859	under J-FGC12, outside J-FGC1695	1350	1855 *	157	5	Husayni 4, Hashimi 1
J-FGC19901	under J-FGC1695, outside J-FGC1721	1450	1500	133	1	Hasani 1
J-Y296615	under J-FGC1695, outside J-FGC1721	750	780	3	1	Hashimi 1
J-ZS5341	under J-FGC1695, outside J-FGC1721	1250	1350	127	1	Hasani 1
J-ZS9151	under J-FGC1695, outside J-FGC1721	1400	—	39	6	Hasani 5, Husayni 1

* Older than the Hashim threshold in Source A, where the clade is therefore not one unit and its claimants lie in subclades that are.
— : not found under that label in Source A. A name is counted once per record, and a record may carry more than one.

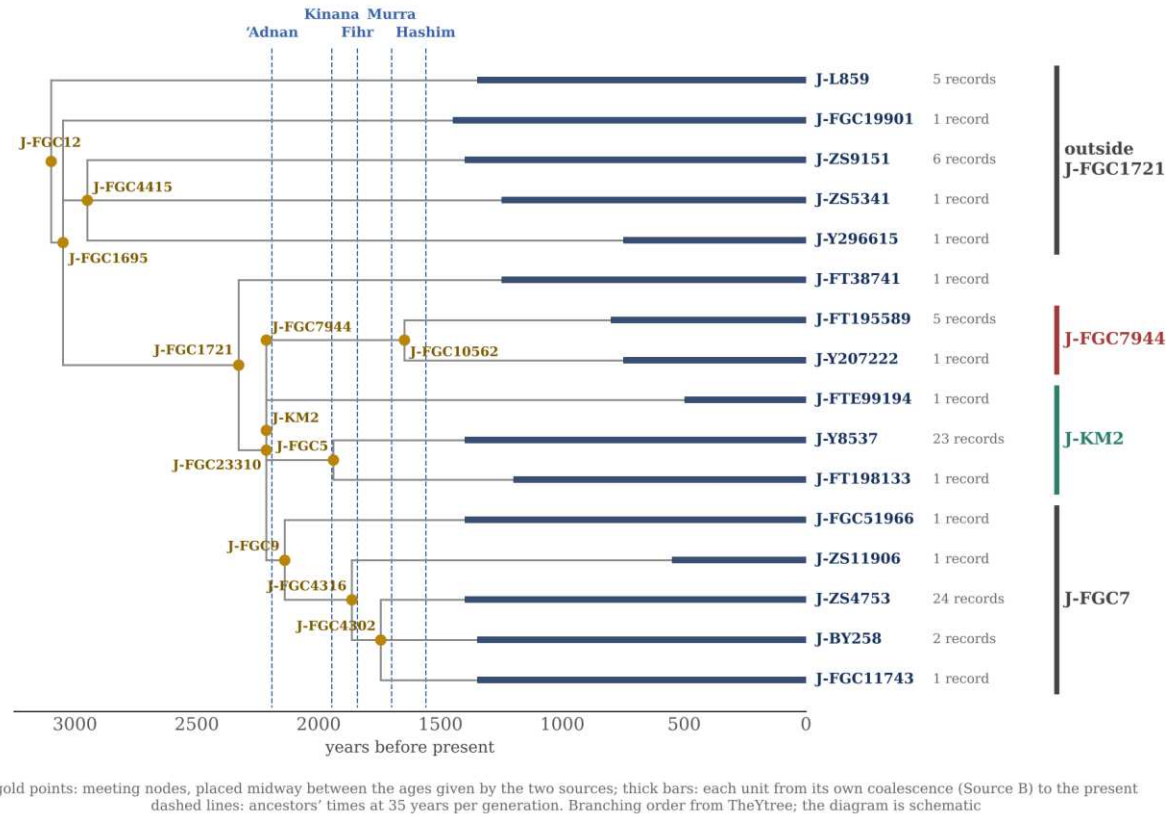


Fig. 1 The sixteen units on the tree. Gold points are the meeting nodes, placed midway between the ages given by the two sources; thick bars run from each unit's own coalescence in Source B to the present; dashed lines mark ancestors' times at the declared ceiling. The brackets name the clades in which the units lie. Branching order follows Source A; the diagram is schematic

The sixteen units hold 75 records. A further 35 claimants in haplogroup J were tested too shallowly to be placed in a unit, and 17 belong to other haplogroups altogether; neither group enters the tables. Section 6.2 and Section 7 follow in detail three meeting nodes, the deepest, the widest and the nearest, through four clades: J-L859; J-FGC7944 (YFull: J-Y10957); and J-ZS4753 and J-FGC4309, both under J-FGC4302 within J-FGC4316 (YFull: J-Y8975). J-BY258 lies within J-FGC4309 and is used in Section 7.6.

The paths of FGC7944 and of J-FGC4316 are parallel. TheYtree codes J-FGC7944 as J1a2a1a2d2b2b2c4d2a2a5a1e1 and J-FGC7, which leads to J-FGC4316, as J1a2a1a2d2b2b2c4d2a2a5a1e2: they are sister branches under J-FGC23310 (equivalent to J-FGC1696). A third branch, J-KM2, stands beside them under the same node and meets both at it. The YFull path to J-FGC4316 likewise runs J-FGC23310 > J-Y9272 (FGC7) > J-Y9274 (FGC9) > J-Y8975, without passing through J-FGC7944. Their meeting node is therefore J-FGC23310, dated 2200 years in Source B and 2230 in Source A. In both databases this node is a polytomy whose branches share a single date.

The Source B export lost the indentation that encodes parentage, so the tree was reconstructed with a stack using that database's own invariant: a child's formation time equals its parent's coalescence time. A node was treated as able to carry descendants only when the export lists a paragroup line (X*) for it. Two regions, each collapsed to a single date, remained ambiguous as to parentage: seven nodes at 1700 years and six at 2200. Any arrangement of parentage within a region gives the common ancestor the same coalescence time,

so no meeting age depends on it. Structural statements do depend on it, and where they matter they are taken from the branch codes of Source A, as above for J-FGC23310.

The collapse originates in Source B's dating grid. Across the 2,072 nodes of the J-L859 and J-FGC23310 subtrees it uses only 46 distinct values, all multiples of 25 years, and the grid coarsens with depth until its upper values are 1500, 1550, 1600, 1650, 1700, 1800, 1900 and 2200. In the region at 1700, 7 nodes share a coalescence time of 1700 and 6 of them a formation time of 1700 as well, which is zero elapsed time between arising and splitting. Source A makes ZS4753 and FGC4309 siblings under J-FGC4302 with a coalescence of 1790, whereas the Source B export leaves undetermined whether ZS4753 lies inside FGC4309 or beside it. No meeting age depends on this.

Apart from the rule of Section 3.6, no name, declaration or classification was taken from the public projects. Samples are not identified and no individual is named; units are given by identifier, age and number of records. A record is one entry in one source, so a person present in both sources is counted twice. The public trees are treated as observational datasets, and no conclusion rests on any single sample.

Table 2 The three meeting nodes followed in detail: ages in the two sources (years before present) and, from Source B, the samples below each node and the SNPs that define it

Meeting node (A / B)	Sides it joins	Source A	Source B	B ÷ A	Samples below	SNPs at node
J-FGC12 / J-ZS2121	L859 with each of the other three	3100	3000	0.97	6,251	18
J-FGC23310	FGC7944 with ZS4753 and with FGC4309	2230	2200	0.99	3,168	1
J-FGC4302 / J-Y10966	ZS4753 with FGC4309	1790	1700	0.95	272	2

The two databases are close on these three ages: the ratio between them runs from 0.95 to 0.99. That closeness is not by itself an argument for robustness, since two trees built in part from the same samples may agree for shared reasons. Section 7 examines how disagreement varies with depth.

Every meeting node stands above hundreds or thousands of samples, so none of the three ages rests on a sparsely sampled node. The first branches below J-FGC23310 and J-FGC4302 are short, between one and three SNPs, which is why those two ages are resolved from full descendant paths in Section 7.5 and not from the first branches alone.

As the property of Section 3.5 requires, each pair of sides meets at a single age in each reconstructed tree. The check confirms the reconstruction rather than adding evidence; its consequence is that the verdicts of Section 6 are verdicts on sides, not on the representatives chosen for them.

5. The ceiling on generation length

The criterion needs a value for generation length in order to turn a minimum generation length into a verdict. That value could be estimated from the tree itself, by assigning particular nodes to particular named ancestors and regressing through them. But that assignment is precisely the external designation which Section 1 rejects.

This paper therefore does not estimate generation length from the tree; it declares a ceiling. The declared ceiling is 35 years per generation. Its source is literature independent of the tree: Fenner (2005)

places the cross-cultural mean between 31 and 32 years, and Wang et al. (2023) place the male mean at 30.7. The ceiling is above both, and a higher ceiling yields fewer exclusions.

A single ceiling does not suit every rung equally. The ceiling is drawn from population means over hundreds of generations, which fits the distant rungs well: the mean of twenty-one generations is close to the population mean and its dispersion is small. At ‘Adnan the ceiling places the ancestor 735 years before the Prophet’s generation, where the published means place him 645 to 672 years before it, so it is generous there. At Hashim it is not. What is measured there is the mean of three fathers only, whose dispersion is far wider, and the threshold of 1561 years corresponds to the year 465 CE, close to the date about 464 CE commonly estimated for Hashim’s birth. At the nearest rung, then, the ceiling reproduces the traditional chronology and leaves it no margin.

Table 3 Ancestral age thresholds $T_0 + k \cdot g$ at five generation lengths (years before 2026, rounded). A pair is excluded at a rung by a central estimate when its meeting age exceeds the threshold

Rung	k	30.7 Wang	32 Fenner	35 ★	45	55
‘Adnan	21	2101	2128	2191	2401	2611
Mudar	18	2009	2032	2086	2266	2446
Ilyas	17	1978	2000	2051	2221	2391
Mudrika	16	1947	1968	2016	2176	2336
Kinana	14	1886	1904	1946	2086	2226
al-Nadr	13	1855	1872	1911	2041	2171
Fihr	11	1794	1808	1841	1951	2061
Ghalib	10	1763	1776	1806	1906	2006
Lu’ayy	9	1732	1744	1771	1861	1951
Ka‘b	8	1702	1712	1736	1816	1896
Murra	7	1671	1680	1701	1771	1841
Qusayy	5	1610	1616	1631	1681	1731
Hashim	3	1548	1552	1561	1591	1621

For the three nodes of Table 2 the central verdicts of Source A are the same at Wang’s mean, Fenner’s mean and the declared ceiling at every rung but one: J-FGC4302, aged 1790 years, is excluded at Ghalib at the published means (thresholds 1763 and 1776) and not at the ceiling (1806). Otherwise they change only at forty-five and above, values no published mean supports (Section 8.1).

6. Results

6.1 The open claim: which units are excluded and which remain

The sixteen units form 120 pairs, and every pair meets at one of nine levels of the tree. Table 4 lists the levels from the deepest. For each it gives the date at which the two sides part and how many years before Hashim that is, both from the central age of Source B, and then the grade under the rule of Section 3.4.

Table 4 The nine meeting levels of the 120 pairs: age in each source, the corresponding date and its distance before Hashim, and the grade (g = 35)

Level	Units that meet there	Pairs	Age (B / A)	Date	Years before Hashim	Grade
J-FGC12, J-FGC1695	L859, FGC19901, Y296615, ZS5341 and ZS9151, with every unit and with one another	65	2900–3000 / 3000–3100	975–875 BCE	1,340–1,440	firm at every rung to ‘Adnan
J-FGC23310	J-FGC7944, J-KM2 and J-FGC7, each against the other two	31	2200 / 2230	175 BCE	640	firm through Ghalib; probable through Kinana
J-FGC1721	FT38741 with those three clades	10	2200 / 2230	175 BCE	640	at least that of J-FGC23310
inside J-KM2	FTE99194 with Y8537 and with FT198133	2	2200 / 2230	175 BCE	640	firm through Murra; probable through Lu’ayy
inside J-FGC7 (J-FGC9)	FGC51966 with the other four units	4	2200 / 2080	175 BCE	640	firm through Ka‘b; probable through Ghalib
inside J-FGC7 (J-FGC4316)	ZS11906 with the three under J-FGC4302	3	1800 / 1930	225 CE	240	firm at Hashim; probable through Murra
inside J-KM2 (J-FGC5)	Y8537 with FT198133	1	1800 / 2080	225 CE	240	undecided at every rung
inside J-FGC7 (J-FGC4302)	ZS4753, BY258 and FGC11743, with one another	3	1700 / 1790	325 CE	140	probable at Hashim; not firm
inside J-FGC7944 (J-FGC10562)	FT195589 with Y207222	1	1650 / —	375 CE	90	undecided at every rung

Date and distance before Hashim: from the central age of Source B and the Hashim threshold of 1561 years, which is 465 CE; the intervals around these ages are in Figure 2 and Table A3. “inside”: a level at which only units of one clade meet.

The first three rows are the verdicts between clades. The five units outside J-FGC1721 part from everything else some fourteen centuries before Hashim. J-FGC7944, J-KM2 and J-FGC7 part from one another more than six centuries before him, about 175 BCE by the central estimates of both sources, and J-FT38741 parts from all three no later. These exclusions do not depend on the fine resolution of any node: the first holds at every rung, and the second is firm through Ghalib under each published rate. The remaining rows concern units inside a single clade. There the meeting is nearer, the interval around it reaches or crosses Hashim’s time, and the grade weakens accordingly.

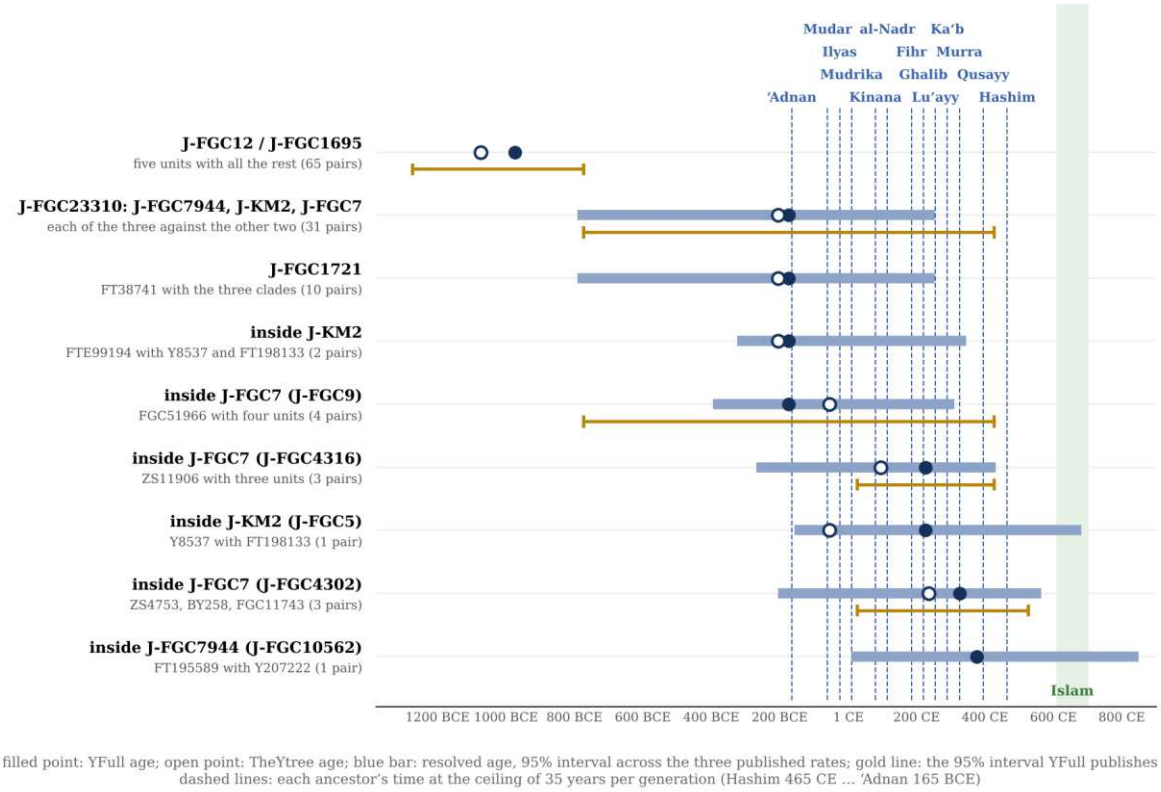


Fig. 2 The meeting levels on a calendar axis. Filled point: Source B age; open point: Source A age; blue bar: the resolved age, 95% interval across the three published rates; gold line: the 95% interval that Source B publishes for the node. Dashed lines mark each ancestor's time at the declared ceiling, from 'Adnan (165 BCE) to Hashim (465 CE); the shaded band is the first century of Islam. A level whose bar lies wholly to the left of an ancestor's line is excluded at that rung

In Source A two of the three pairs at J-FGC4302, those of BY258, meet at that node itself; the third, ZS4753 with FGC11743, meets one node lower, at J-KHU31, which Source A dates to the same 1790 years. The grades of J-FGC4302, J-FGC23310, J-FGC4316 and J-FGC9 are computed in Section 7.5 (Table 6), and those of the three levels inside J-FGC7944 and J-KM2 in Table A4. J-FGC1721 is an ancestor of J-FGC23310, so the pairs that meet there are excluded at least as far.

Source B publishes its own interval for five of these nodes, the gold lines of Figure 2, and they are wider than the resolved ones. At Source B's own rate the lower end of the resolved interval stands 57 and 97 years above the lower end of Source B's interval for J-FGC4302 and J-FGC4316, and 228 and 288 years above it for J-FGC9 and J-FGC23310 (Table A3). The grades therefore claim more precision than Source B does, most of all at the two deeper nodes. Read on Source B's intervals alone, the deep level is excluded at every rung, J-FGC23310, J-FGC9 and J-FGC4316 at Hashim only, and J-FGC4302 at no rung.

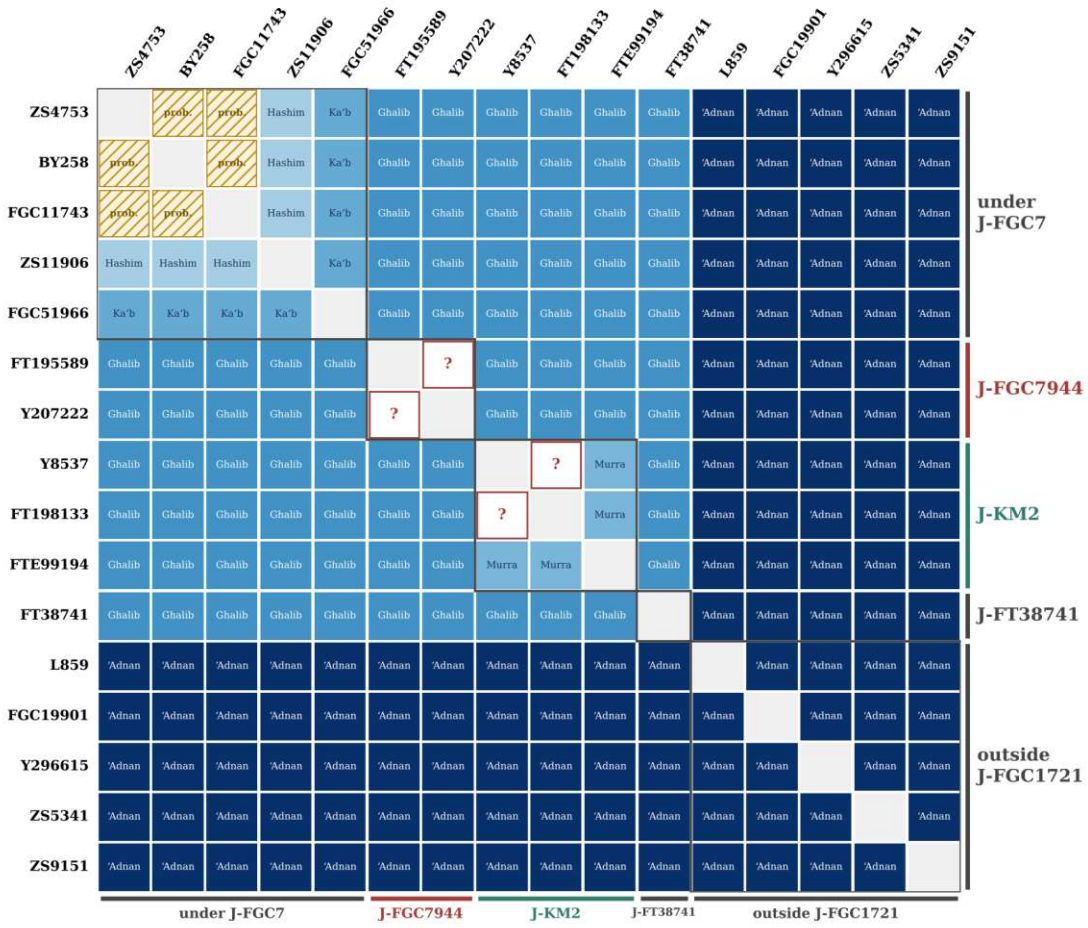


Fig. 3 The verdict on every pair of units. Each cell names the highest ancestor at which joint descent of the pair is firmly excluded; ‘Adnan means every rung. Hatched cells are excluded at Hashim with a probability of 0.80–0.95 only, and cells marked “?” are undecided at Hashim. The blocks along the diagonal are the clades, and every cell outside the blocks is a verdict between clades. g = 35

Figure 3 gives the same result pair by pair. Read as an order of separation, descending from ‘Adnan: the five units outside J-FGC1721 (14 records) stand apart from everything else and from one another at every rung. J-FT38741, the two units of J-FGC7944 (6 records) and the three of J-KM2 (25 records) then part from one another and from J-FGC7 through Ghalib. Inside J-FGC7, J-FGC51966 leaves through Ka‘b and J-ZS11906 at Hashim. The three that remain to the last, J-ZS4753, J-BY258 and J-FGC11743 (27 records), are probably but not firmly separated at Hashim.

The order is an order of separation and nothing more. A unit that leaves early is not thereby shown to carry a false claim, and the units that remain to the last are not thereby shown to carry a true one. Exclusion is a property of pairs, and what the figure establishes is a count. Of the 120 pairs, 115 are firmly excluded from joint descent from Hashim, three are probably excluded and two are undecided. The firm exclusions divide the sixteen units into twelve sets, the three under J-FGC4302, the two in J-FGC7944, J-Y8537 with J-FT198133, and nine single units, such that no two sets can both descend from Hashim; within a set the

test does not decide. The number of records plays no part in any verdict; it is reported because the rule that fixes the sides rests on what members wrote.

6.2 The three levels followed in detail

Three of the levels are followed more closely here and in Section 7, through the four clades named in Section 4: J-FGC12, where L859 meets the rest; J-FGC23310, where FGC7944 meets J-FGC4316; and J-FGC4302, where ZS4753 meets FGC4309.

Whichever side holds the measure, L859 is excluded with every other side at every rung, even at ‘Adnan, the least demanding: joining there would require 78.3 years per generation in Source A and 73.5 in Source B.

FGC7944 is excluded with ZS4753 and with FGC4309 at every rung by the central estimates of both databases, but the margins above Kinana are small. Under the three published rates the exclusion is firm through Ghalib, probable from Fihr to Kinana, and undecided from Mudrika to ‘Adnan (Table 6).

ZS4753 and FGC4309 are excluded with each other at Hashim by the central estimates of both databases, which the definition of a unit guarantees for Source B. The probability that their divergence predates Hashim is 0.86 under the fastest published rate and 0.97–1.00 under the other two, so the exclusion is probable and not firm (Table 6). At Qusayy and above the instrument is silent.

Table 5 Minimum generation length for joint descent through each meeting node (years per generation); values above 35 are excluded at the declared ceiling

Meeting node	M	‘Adnan	Ilyas	Mudrika	Kinana	Fihr	Murra	Hashim
J-FGC12 (A)	3100	78.3	96.7	102.8	117.4	149.5	234.9	548.0
J-ZS2121 (B)	3000	73.5	90.8	96.5	110.3	140.4	220.6	514.7
J-FGC23310 (A)	2230	36.9	45.5	48.4	55.3	70.4	110.6	258.0
J-FGC23310 (B)	2200	35.4	43.8	46.5	53.1	67.6	106.3	248.0
J-FGC4302 (A)	1790	15.9	19.6	20.9	23.9	30.4	47.7	111.3
J-Y10966 (B)	1700	11.6	14.4	15.3	17.4	22.2	34.9	81.3

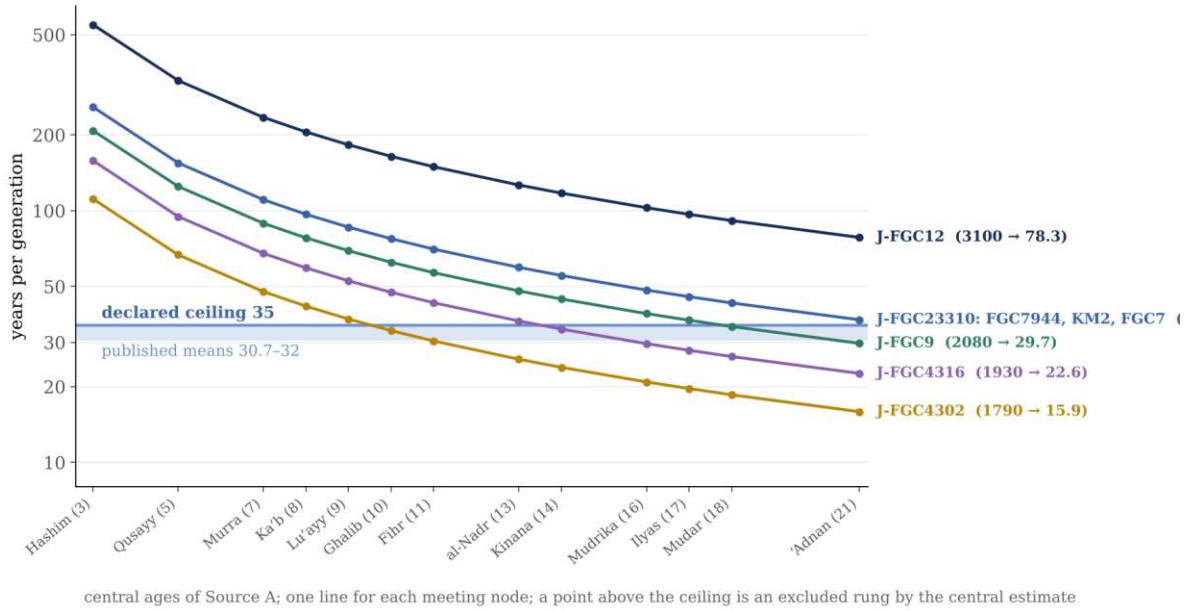


Fig. 4 Minimum generation length for joint descent through each meeting node (Source A). The vertical axis is logarithmic; the band runs from the published means to the declared ceiling, and a line above the ceiling marks an excluded rung

A verdict on a meeting node applies to every lineage below it. Any subclade of FGC7944, however young, meets any subclade of FGC4316 at J-FGC23310, and any subclade of L859 meets any lineage of the other sides at J-FGC12. The youth of a subclade does not bring its meeting node closer.

The phrase "not excluded" is deliberate. The criterion refutes and does not confirm: that a joining is not chronologically excluded is not a joining, and is no evidence that the two sides descend from that ancestor.

Nor does non-exclusion at a rung place the meeting at that rung. The instrument gives a lower bound only: it finds the highest rung below which joint descent is impossible, and it does not locate the meeting above that rung. The actual meeting may lie far above the first rung that is not excluded. That FGC7944 with FGC4316 is undecided from Mudrika upward does not mean that the two meet at Mudrika; locating the meeting would require a mapping of the tree onto the genealogy, which lies outside this paper.

7. Stability

7.1 The observed distribution

The two databases were compared on 880 nodes shared by name within the sides studied. The ratio of the B estimate to the A estimate has median 1.00; the median of |1 − ratio| is 14.1% and its mean 26.2%. B estimated the smaller value in 411 nodes and the larger in 295, and the two agreed exactly in 174.

This distribution describes two particular databases; it does not measure uncertainty. Their agreement bounds the error in neither direction: had they agreed on one figure and been wrong together, the agreement would appear perfect and the error be large. What follows is an observed rate between two databases, read as a description of how stable a verdict is when the database is changed. An independent check is provided by the STR clock of Section 7.6.

7.2 Disagreement and node depth

Disagreement falls from 21.4% at the shallowest nodes to 8.7% at 1200–1600 years, then rises to 14.2% in the deepest band. The deepest band holds only 49 nodes, and across the whole J-FGC11 tree (1,913 shared nodes) the same band shows 7.9%, so the rise is specific to the sides studied. A shallow node's age is estimated from few mutations in few samples, whereas a deep node is supported by many.

The three meeting nodes lie between 1700 and 3100 years, entirely within the deepest band. This is not a choice made to favour the result: the one-sided constraint requires it before any data are examined. Since the deepest band does not show the lowest disagreement here, the directional test of 7.3 and the methods of 7.5 and 7.6 are applied.

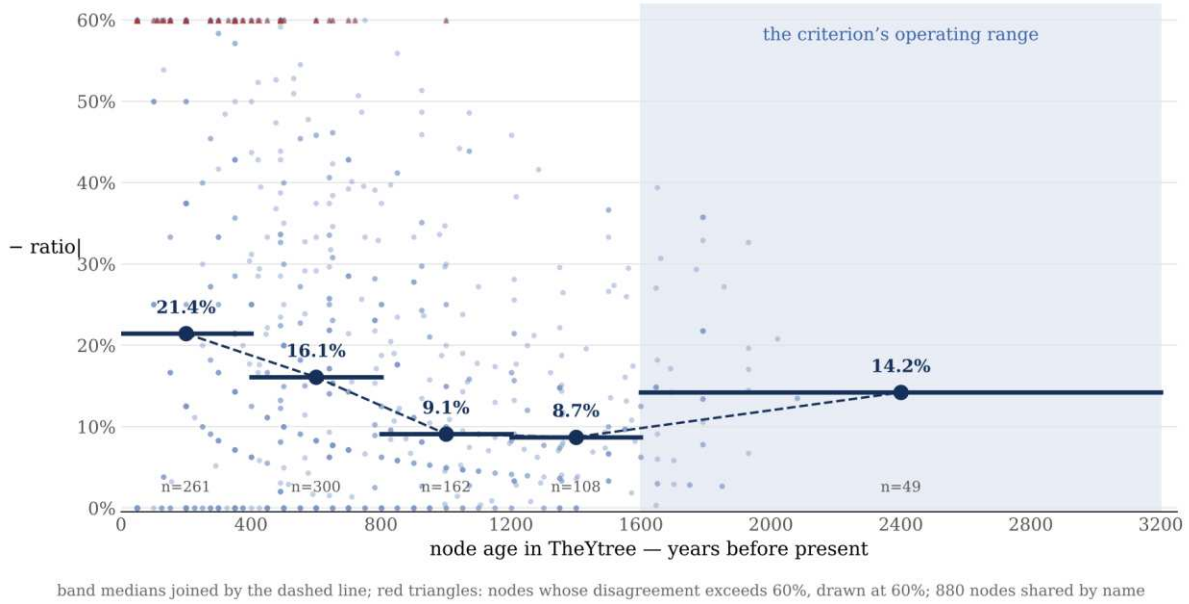


Fig. 5 Disagreement between the databases by node age, 880 shared nodes. Each point is a shared node and each dark bar the band median; nodes whose disagreement exceeds 60% are drawn as triangles at 60%; the blue region is the criterion's operating range

7.3 The direction of disagreement in the deep range

In the shallow range disagreement is balanced: the median ratio is 1.000, and Source B is lower in 294 of 723 nodes. In the deep range it is directional: Source B is lower than Source A in 117 of 157 nodes aged 1200 years or more, with a median ratio of 0.909, a systematic shrinkage of 9.1%.

A frequency cannot see such a bias, since it moves every node the same way at once. It is therefore tested directly: what survives of the Source A verdicts if every one of its ages is reduced by 9.1%? The answer, over the thirteen rungs, is that seven central verdicts are overturned and no more: J-FGC23310 at 'Adnan, Mudar and Ilyas, and J-FGC4302 at Qusayy, Murra, Ka'b and Lu'ayy (1790 years reduced by 9.1% is 1627, below the Qusayy threshold of 1631). Every central verdict on J-FGC12 survives at every rung, and every one at Hashim survives. The distance of J-FGC12 from the ceilings ranges from 29% to 50%, which no shrinkage of this size approaches.

This test does not exhaust systematic error: an error common to both databases cannot appear in it. Its use is to separate the verdicts that a bias of this size overturns from those it does not.

7.4 The stability of a verdict between the databases

Of every exclusion two questions are asked: what shrinkage of the meeting age would overturn it, and how often has a disagreement that large occurred between the two databases? A verdict computed on Source A is overturned if A is too high, and the evidence for that is the nodes where B is lower than A; a verdict computed on Source B is overturned if B is too high, and the evidence is the nodes where A is lower than B. Each column therefore has its own tail.

Three bands are adopted: high where the frequency is twelve per cent or less, moderate where it exceeds twelve and does not exceed twenty-five, and low where it exceeds twenty-five. They are presentational cut-offs, not natural boundaries, and they measure neither the truth of a claim nor its probability. The frequency is given on all 880 nodes and, in parentheses, on the 157 nodes aged 1200 years or more, where the criterion operates; the deep sample is small, so no band is built on it.

The exclusion at J-FGC12 is of high stability at every rung in both sources, except at ‘Adnan in Source A, where it is moderate. The exclusion at J-FGC23310 is of moderate or high stability from Hashim to Fihr in both sources, and above Fihr of low stability in Source A and moderate in Source B up to Mudrika, which is why Section 7.5 gives its probabilities directly. The exclusion at J-FGC4302 is of low stability on all nodes; on deep nodes the Source B frequency at Hashim falls to 4%, and Section 7.5 again gives the probability directly.

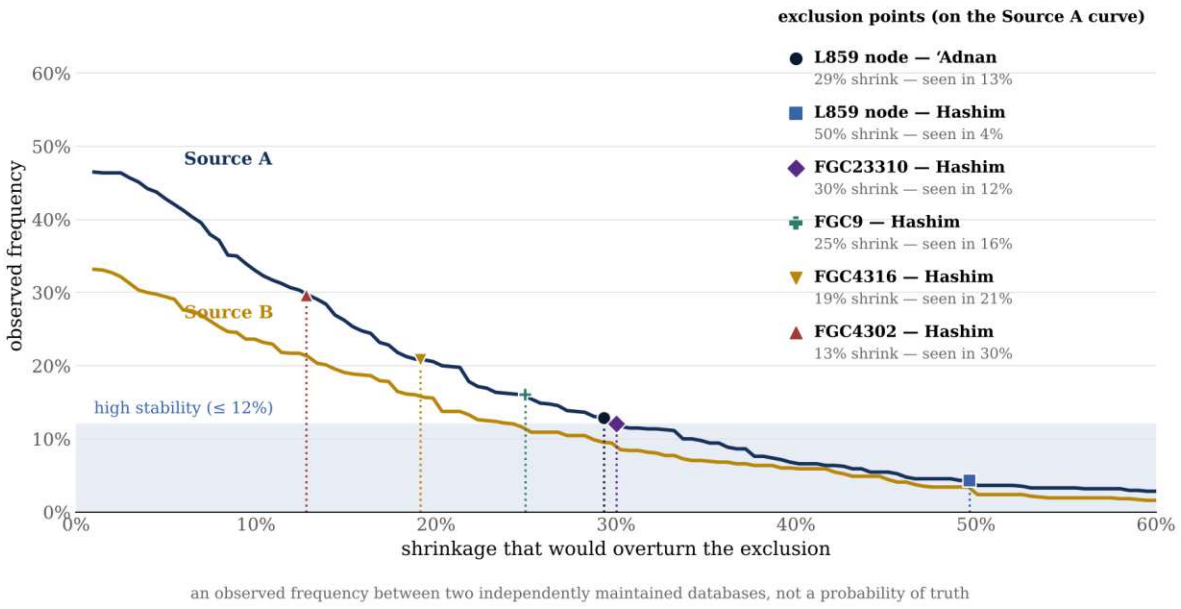


Fig. 6 Stability of exclusion. Each source has its own tail in the direction that would overturn its verdicts; the points lie on the Source A tail. The blue region is the high-stability zone, where the frequency is twelve per cent or less

7.5 Branch SNP counts and the mutation rate

Source B assigns several distinct divergences the same age, because its dating grid coarsens with depth. The branches between those nodes nevertheless carry SNPs, and SNPs require time. For a node X and a

well-dated descendant node D, X is older than D by the time needed to accumulate the n SNPs on the path between them:

$$T_x = T_D + \Delta_n, \quad \Delta_n \sim \text{Gamma}(n + a, r)$$

Here r is the number of years per listed SNP and a is a prior offset, set to zero in the base case; a = 0.5 and a = 1 are reported in Table A2. Across all of Source B's branches, those of zero length included, the sum of branch durations divided by the sum of listed SNPs is 81 years per SNP; over the branches of positive length alone it is 89. The model uses 80 with a standard deviation of 6, at or below both figures, and a lower figure gives younger meeting ages and so fewer exclusions. Uncertainty in T_D is modelled as a relative error of 10–50%, according to the number of independent lineages below D. These errors are twice those first assigned by judgement. They were widened so that the resolved interval for J-FGC4302 at Source B's own rate, 1557–2072 years, is no narrower than the interval Source B publishes for that node, 1500–2000 (Table A3). Each node is estimated from several descendant paths, combined by inverse-variance weighting, and the agreement between paths is tested with Cochran's Q.

The calibration is the decisive premise. An error in the mutation rate moves every age at once, the descendant ages included, so it cannot be detected by comparing two databases or two paths. It is therefore modelled directly: all ages, and the years per SNP, are multiplied by the ratio of Source B's rate to an alternative rate. Three published rates are used. Source B states that it uses 0.8178×10^{-9} per base pair per year, calibrated on the ancient Anzick-1 genome and on known genealogies, with a 95% interval of about $0.70–0.94 \times 10^{-9}$ (Adamov et al. 2015; YFull 2026b). Fu et al. (2014), from ancient genomes, give 0.76×10^{-9}, which would make the ages about 8% older; Helgason et al. (2015), from deep Icelandic pedigrees, give 0.871×10^{-9}, which would make them about 6% younger. The upper end of Source B's own interval, 0.94×10^{-9}, would make them about 13% younger, and is reported as a limit rather than used in the rule.

Do the ages depend on a generation length? The test sets an age in years against a ceiling on generation length, and it would be circular if the ages had themselves been obtained by assuming a generation length. They were not. Fu et al. (2014) calibrate on a radiocarbon-dated ancient genome, counting the mutations it lacks by reason of its age, which yields a rate per year with no generation length in it. Helgason et al. (2015) count mutations between Icelandic men whose common ancestors have recorded dates of birth, so their rate is likewise per year. Source B anchors its rate on the radiocarbon-dated Anzick-1 genome together with documented genealogies (Adamov et al. 2015). None of the three rates converts generations into years, and the rule of Section 3.4 requires all three.

For J-FGC4302 (YFull: J-Y10966) the seven paths give Q = 1.9 on 6 degrees of freedom (p = 0.93); the median age is 1802 years under Source B's rate, 1939 under Fu's and 1692 under Helgason's. For J-FGC23310 the four paths give Q = 2.5 on 3 degrees of freedom (p = 0.47), and the median age is 2238, 2408 and 2101 years respectively. The method is not independent of Source B: it uses its SNP calls and its ages for the descendant nodes. What it adds is the information on the stems that the dating grid discards.

The open claim of Section 6.1 brings in two further meeting nodes, each lying between J-FGC4302 and J-FGC23310, and the same computation is applied to them. J-FGC4316 (YFull: J-Y8975) joins the three units under J-FGC4302 to the unit J-ZS11906: fourteen paths give Q = 7.2 on 13 degrees of freedom (p = 0.89) and a median age of 1908 years under Source B's rate, 2053 under Fu's and 1791 under Helgason's. J-FGC9 (YFull: J-Y9274) joins the unit J-FGC51966 to everything under J-FGC4316: seventeen paths give Q = 14.2 on 16 degrees of freedom (p = 0.58) and a median age of 2033, 2187 and 1908 years. The four

resolved ages fall in the order the tree requires, J-FGC4302 the youngest and J-FGC23310 the oldest, although nothing in the computation imposes that order. Three small levels inside J-FGC7944 and J-KM2 are resolved in the same way in Table A4.

Table 6 Probability that the divergence at each of the four resolved nodes predates each ancestor (g = 35): the range over the three published mutation rates, with the grade under the rule of Section 3.4

Ancestor	k	Threshold	J-FGC4302	J-FGC4316	J-FGC9	J-FGC23310
Hashim	3	1561	0.86–1.00 ○	0.99–1.00 ●	1.00 ●	1.00 ●
Qusayy	5	1631	0.69–0.99	0.94–1.00 ○	1.00 ●	1.00 ●
Murra	7	1701	0.47–0.96	0.81–1.00 ○	0.98–1.00 ●	0.99–1.00 ●
Ka‘b	8	1736	0.36–0.93	0.70–1.00	0.96–1.00 ●	0.99–1.00 ●
Lu’ayy	9	1771	0.26–0.89	0.58–0.99	0.92–1.00 ○	0.98–1.00 ●
Ghalib	10	1806	0.18–0.83	0.44–0.98	0.85–1.00 ○	0.96–1.00 ●
Fihr	11	1841	0.12–0.76	0.32–0.97	0.75–1.00	0.94–1.00 ○
al-Nadr	13	1911	0.04–0.58	0.13–0.89	0.49–0.99	0.87–1.00 ○
Kinana	14	1946	0.02–0.48	0.07–0.82	0.36–0.98	0.82–0.99 ○
Mudrika	16	2016	0.01–0.29	0.02–0.62	0.15–0.94	0.69–0.98
Ilyas	17	2051	0.00–0.22	0.01–0.51	0.08–0.89	0.61–0.97
Mudar	18	2086	0.00–0.15	0.00–0.39	0.04–0.81	0.53–0.95
‘Adnan	21	2191	0.00–0.04	0.00–0.13	0.00–0.49	0.30–0.87

● firm: at least 0.95 under each rate; ○ probable: at least 0.80 under each; unmarked: undecided. The lowest figure in every cell is that of the fastest rate (Helgason, 0.871) and the highest that of the slowest (Fu, 0.76). J-FGC4302: ZS4753, BY258 and FGC11743 with one another (seven paths); J-FGC4316: ZS11906 with those three (fourteen); J-FGC9: FGC51966 with those four (seventeen); J-FGC23310: J-FGC7944, J-KM2 and J-FGC7 with one another (four). Probabilities at the upper end of Source B’s rate interval, 0.94, are in Figure 7.

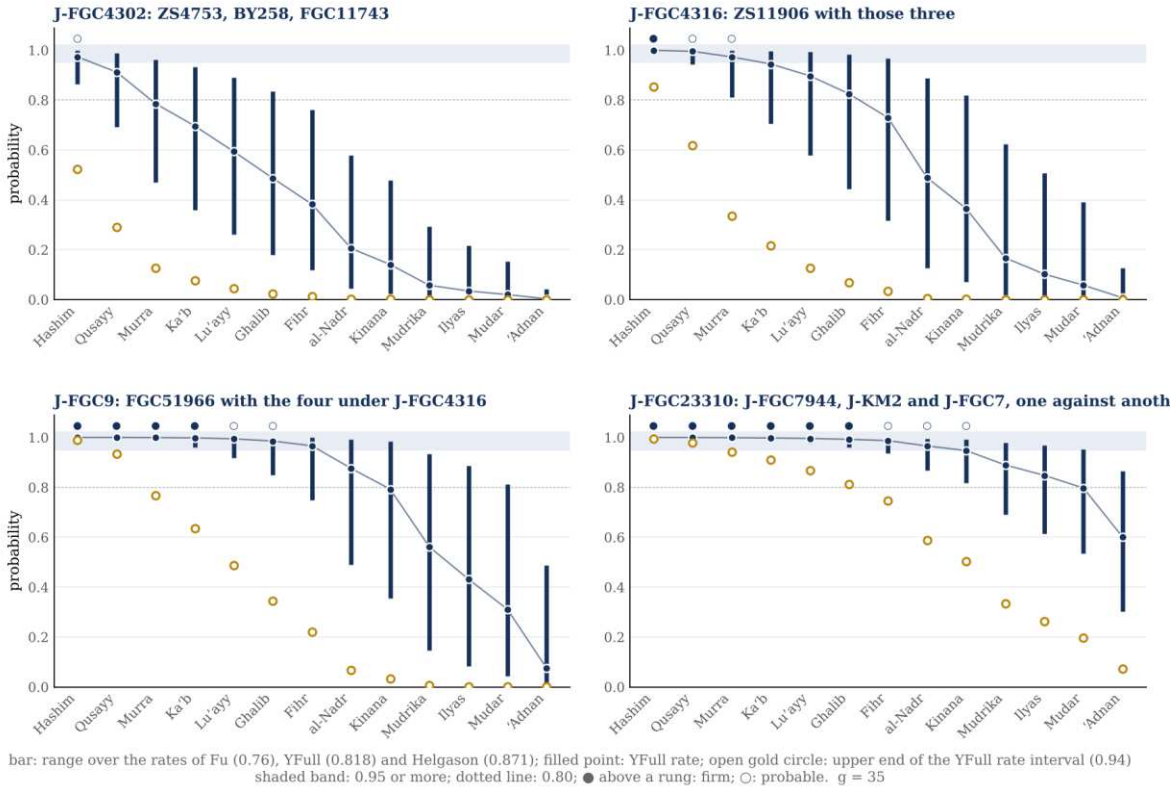


Fig. 7 Probability that each divergence predates each ancestor, from branch SNP counts (g = 35), for the four resolved nodes. Each bar spans the three published rates; the point marks Source B's own rate, and the open gold circle the upper end of its rate interval. The shaded band marks probabilities of 0.95 or more; ● above a rung marks a firm exclusion and ○ a probable one

The probabilities are conditional on the model as well as on the rate. Correlation between paths that share an upper branch is ignored, which makes the intervals too narrow; the widened errors on T_D are meant to offset this and do not model it. Values are given to two decimals only.

7.6 A third, independent clock from STR data

STR markers offer a clock whose measurements depend on neither database: their mutations are counted on different loci, by different laboratories, with no use of Source A or Source B ages for the between-group distances. Its conversion to years, however, is calibrated on Source B (below), so it is independent in measurement and not in calibration.

Data. STR results were taken from the public results pages of three FamilyTreeDNA group projects. Kits were deduplicated and each was placed on Source B by its SNP-confirmed haplogroup; kits with predicted haplogroups were excluded. The analysis uses 93 single-copy markers present in at least 90% of kits. Multi-copy markers (DYS385, DYS459, DYS464, YCAII, CDY, DYF395S1, DYF406S1, DYS413, DYS425) were excluded, DYS389ii was replaced by DYS389ii − DYS389i, and null alleles (recorded as 0) were treated as missing. Sample sizes: L859 83, FGC7944 (FGC10579 branch) 36, ZS4753 44, FGC10092 34, FGC4320 35, BY258 12.

Method. The distance between two groups is the average squared difference (ASD) over all cross-group pairs, summed over markers. It is converted to years by calibration on within-group ASD against the mean

pairwise TMRCA of the same group in Source B; each of the six groups gives one calibration. Large single-event jumps inflate ASD without reflecting time, so differences were also capped at two steps and at one step per marker.

Table 7 STR estimates of divergence age (years), range over the six calibrations

Pair	Source B	Uncapped	Cap 2 steps	Cap 1 step
L859 × ZS4753	3000	2140–3430	2530–3130	2330–2980
L859 × BY258	3000	2110–3380	2440–3020	2250–2880
L859 × FGC7944	3000	2450–3920	2720–3370	2330–2980
FGC7944 × ZS4753	2200	1170–1880	1480–1830	1420–1820
FGC7944 × BY258	2200	1250–2010	1550–1910	1510–1930
FGC7944 × FGC4320	2200	1560–2490	1900–2350	1520–1940
ZS4753 × BY258	1700	1060–1700	1360–1680	1360–1730

Within groups, STR ages are close to Source B: expressed as ratios to the ZS4753 × FGC4309 divergence, Y92662 gives 0.74 against 0.74 and FGC4320 0.79 against 0.71. Between groups, STR ages are mostly younger than the SNP estimates of Section 7.5 (for J-FGC4302, about 1360–1680 against a median of 1692–1939), the expected effect of saturation in fast markers, and they vary with the group chosen: the three pairs across J-FGC23310 give 1480–1830, 1550–1910 and 1900–2350 with a cap of two steps. The STR clock is therefore conservative at depth: an exclusion it confirms is robust, and an exclusion it fails to confirm is not refuted.

The uncapped L859 calibration was an outlier because of rare large jumps (DYS442 = 7 in five kits across different L859 sub-branches, modal 12; DYS441 = 10 in two kits of one sub-branch, modal 16). Capping removes the outlier, and all six calibrations then agree. With any capping, every estimate for L859 lies above the ‘Adnan threshold (2191). Without capping, the lowest of the six calibrations for L859 × BY258 (2110) lies above the Mudar threshold (2086) and below that of ‘Adnan. The third clock therefore confirms the exclusion at J-FGC12 through Mudar without condition, and at ‘Adnan on condition that the jumps are capped.

8. Sensitivity

8.1 The generation-length threshold

Table 3 gives the thresholds at five generation lengths. For the three nodes of Table 2 the central verdicts of Source A are the same at 30.7, 32 and 35, except at Ghalib for J-FGC4302 (Section 5). At forty-five, J-FGC23310 ceases to be excluded at ‘Adnan and Mudar; at fifty-five it ceases to be excluded from Mudrika upward, and J-FGC4302 ceases to be excluded at Murra (1841 against 1790). The central exclusion at J-FGC12 holds at all five values.

The narrowest margins lie at Murra and at ‘Adnan. At Murra the threshold of 1701 years stands one year above the Source B meeting age of 1700, so the central estimate there is indeterminate; at the published means the threshold falls to 1671 and 1680 and the node is excluded in both sources. At ‘Adnan the meeting age at J-FGC23310 exceeds the threshold by 9 years in Source B and 39 in Source A. Source B’s displayed ages at this depth move in steps of 100 years or more (Section 4), so a margin of 9 years is below its

resolution. These margins are why the verdicts are given as probabilities in Section 7.5. At Hashim the central margins are 229 years at J-FGC4302 in Source A and 139 in Source B.

8.2 The ancestor to whom the name Quraysh attaches

Two views are held as to which ancestor the name Quraysh attaches to: Fihr b. Malik, giving k = 11, and al-Nadr b. Kinana, giving k = 13. At the declared ceiling their thresholds are 1841 and 1911. Both lie below the meeting ages of J-FGC12 and J-FGC23310, so those nodes are excluded on either view, and both lie above that of J-FGC4302 (1790), so neither view adds an exclusion there. The verdicts are the same whichever ancestor the name is attached to.

8.3 The reference year for "present"

The reference year used for ages in "years before present" must be consistent with the historical calibration point. YFull defines "ybp" relative to the current year rather than to 1950 (YFull 2026c), so for the version used, dated 2 November 2026, the main calculations use 2026 and $T_0 = 2026 - 570 = 1456$ years. The convention of TheYtree was not established from its documentation; its ages are treated as relative to 2026, which is an explicit working assumption for Source A.

To examine that assumption, a hypothetical 1950 reference was applied to Source A, with $T_0 = 1380$ and the exported ages held fixed. Changing T_0 while holding M fixed raises the minimum generation length by $76 \div k$: about 25.3 years per generation at Hashim and 3.6 at ʻAdnan. For J-FGC12 the minimum rises from 548.0 to 573.3 at Hashim and from 78.3 to 81.9 at ʻAdnan; for J-FGC23310 from 258.0 to 283.3 and from 36.9 to 40.5; for J-FGC4302 from 111.3 to 136.7 and from 15.9 to 19.5. The alternative reference therefore makes the test more restrictive.

At the declared ceiling two central verdicts in Source A change, both in the direction of exclusion: the threshold at Fihr falls from 1841 to 1765, below the meeting age of J-FGC4302 (1790), so that node becomes excluded at Fihr; at Ghalib the threshold falls to 1730 and the same holds. At al-Nadr the threshold is 1835 and the node remains not excluded. No verdict at J-FGC12 or J-FGC23310 changes. The Source B calculations retain the documented 2026 reference.

8.4 The generation count k

The generation count is taken from the attested genealogy, which is massively transmitted for the rungs below ʻAdnan. Should a rung or two above Mudrika be disputed, the effect on the denominator is slight: adding one generation to twenty-one reduces the minimum generation length at J-FGC12 from 78.3 to 74.7 years at ʻAdnan, both far beyond the plausible. At the lower rungs a single generation matters more, which is why Table 5 reports the rungs separately; it shows seven of the thirteen, and Table 3 gives the thresholds of all.

8.5 The mutation rate

Because a rate error scales every age by the same factor, the shrinkage column of Table A6 is also the proportion by which the published ages would have to be too old for each verdict to fall. Helgason's rate corresponds to ages about 6% too old and the upper end of Source B's interval to about 13%. The node at which L859 meets the other sides would fall only if the ages were too old by 27–50%, according to the rung, which no published rate approaches. J-FGC23310 falls within Helgason's rate at ʻAdnan (0.4–2%)

and, in Source B, at Mudar (5%), and within the limit of 13% at Mudar in Source A and at Kinana, Mudrika and Ilyas (6–13%). J-FGC4302 falls at Murra within Helgason's rate (0–5%) and at Hashim within the limit (8–13%). TheYtree does not state its rate, so the same scaling is assumed for Source A. These are the reasons the rule of Section 3.4 rests on the three published rates, and the grades of Table 6 should be read with this dependence in view.

Table A7 gives, for each meeting node and rung, the fastest rate at which the exclusion remains firm. An exclusion is firm under the published rates when the figure is 0.87 or more, and the distance of the figure from 0.87 is the margin of the verdict. The second figure in each cell is for a generation of 30 years, near the published means: a shorter generation lowers every threshold and so raises every figure, which is why the ceiling of 35 years is the hard case. No single statement covers every node. The exclusions at J-FGC12 hold at any rate up to 0.98; those at J-FGC23310 hold through Qusayy up to 0.97 and through Ghalib up to 0.87; that at J-FGC4316 holds at Hashim up to 0.90; and that at J-FGC4302 is firm at Hashim only at rates of 0.83 or less, which leaves out the fastest of the three published rates.

9. Limitations

The two databases are maintained independently but may share sample data, which may conceal errors common to both; their observed disagreement does not quantify total dating uncertainty, and nothing in Table A6 is a bound on it in either direction. The STR clock is independent of both, but it is conservative at depth and can only confirm an exclusion; at ʻAdnan its confirmation depends on capping.

Disagreement in the deep range is directional. Its effect was tested by an across-the-board shrinkage, which overturned seven central verdicts while every one at J-FGC12 and every one at Hashim survived. That test cannot show what the two databases share.

The deep verdicts depend on the calibration of the mutation rate, which moves every age at once and cannot be seen in any comparison between databases or paths. The grades rest on three published rates; under the upper end of Source B's own rate interval only the exclusions at J-FGC12 at every rung, at J-FGC23310 at Hashim and Qusayy and at J-FGC9 at Hashim remain at a probability of 0.95 or more; at Fihr that at J-FGC23310 falls to 0.75. The ceiling is declared, not estimated, and its source is literature on general populations rather than on these lineages; the effect of five thresholds is shown so the reader may judge. TheYtree states neither its present-reference, its mutation rate, nor its calibration method, and publishes no version number; YFull states all three (YFull 2026b, 2026c).

The test gives no weight to non-exclusion, and no method that uses the tree alone can give it weight. Two claims lying close together in the tree remain compatible down to the age of their divergence whatever their truth; deciding between them requires an anchor from outside the tree, and the paper uses none. For the same reason it does not identify which side of an excluded pair fails.

The Source B tree was reconstructed from a flattened export, and two regions collapsed to single dates remained ambiguous as to parentage. No meeting age depends on them; structural statements in them are taken from the branch codes of Source A. At Murra the nominal clearance in Source B lies within its dating resolution, so the central-value verdict there is indeterminate and the probability of Section 7.5 decides it.

The SNP-count method is not independent of Source B: it uses Source B's SNP calls and its ages for the descendant nodes, so a bias shared by all of them would pass into it, and correlation between paths was ignored. Three of the seven paths below J-FGC4302 were counted on the flat arrangement of the collapsed

region; the nested arrangement of Source A would add SNPs to them and raise the probabilities, so the grades at that node are conservative. The meeting of FGC7944 and J-FGC4316 lies in a polytomy that both databases date with a single value, and its age rests on four paths whose estimates spread widely, though not significantly. The STR clock is independent of both databases in measurement but is calibrated on Source B.

Source B publishes its own 95% intervals for the meeting nodes, and they are wider than the model's: 1500–2000 years for J-FGC4302, 1600–2000 for J-FGC4316, 1600–2800 for J-FGC9 and for J-FGC23310, and 2800–3300 for J-ZS2121 (Table A3). An interval for one node taken alone is expected to be wider than an estimate that combines several paths to it, but the paths are not independent and the model's precision cannot be checked from outside. A reader who prefers Source B's intervals should read the results thus: the deep level is excluded at every rung; J-FGC23310, J-FGC9 and J-FGC4316 at Hashim only; and J-FGC4302 at no rung.

The rule that fixes the sides reads what members of public projects wrote. A lineage whose members wrote none of the four names does not enter, whatever is claimed for it elsewhere, and a record enters whenever one of the four words appears in it, whether or not it asserts descent. No record was dropped. In J-Y8537, 21 of the 23 records carry Hashimi as a family name; in J-ZS9151 the six records attach Hasani or Husayni to a named tribe of Madhhij; and in J-FGC11743 and J-FGC19901 the single record carries the word Hasani with nothing to show what it means. A reader who discounts any of these removes the corresponding unit and its pairs, and no verdict on any other pair changes. A further 35 claimants in haplogroup J were tested too shallowly to be placed. Units were defined in Source B, which holds most of the samples, so that two of them, J-L859 and J-Y8537, are older than the Hashim threshold in Source A, where their claimants fall in subclades. And because a unit is defined by the Hashim threshold in Source B, exclusion at Hashim by that database's central estimates is built into the list of units and is not a result (Section 3.6).

Two single-record units lie close to larger ones: J-Y207222 beside J-FT195589 in J-FGC7944, and J-FT198133 beside J-Y8537 in J-KM2. Their pairs with their neighbours are the least resolved in the paper: they rest on six and four paths and are undecided at Hashim (Table A4). A third, J-FTE99194, meets the other two units of J-KM2 at J-KM2 itself. That node was first resolved from its two largest immediate branches, which left the exclusion just short of firm at Hashim (0.94). It was then resolved in the manner of the four main nodes, from the dated nodes below its branches, thirteen paths chosen by the rule stated under Table A1, and the grade reported is from those. The rule was fixed after the first result had been seen. The parentage of J-FTE99194 is ambiguous in the Source B export, within the region collapsed at 2200 years, and its pairs are graded at J-KM2, where Source A places it.

The grades depend also on the errors assigned to the ages of the descendant nodes and on the prior offset. Table A2 in the appendix shows what happens when the errors are halved, that is returned to the values first assigned, or raised by half again, when the offset is 0.5 or 1, when the weighting is dropped, and when each path is removed in turn. The grades move by up to three rungs. At J-FGC4302 the exclusion at Hashim becomes firm with the errors halved or with a = 1, and with the errors raised by half it is no longer even probable. The verdict on the three pairs under that node therefore rests on the precision assigned to those ages more than on anything else.

The rates of Fu et al. and of Helgason et al. were applied to Source B's ages by simple proportion. Published rates differ in part because they are measured over different callable regions of the chromosome (Balanovsky 2017), so the proportion is an approximation, and the three rates are better read as a spread of

plausible calibrations than as three exact alternatives. The date of about 464 CE for Hashim's birth, used in Section 5 to judge the ceiling, is a common estimate and not a transmitted date.

Two things must be kept apart. The rule confines the list of sides to lineages whose members declared a claim in the public record: a lineage that declared nothing was not tested, and nothing is said of it. The verdict on a pair of clades is not so confined. A meeting age belongs to the tree and not to the names, so an exclusion holds for every carrier of the two clades, declared or not. The paper therefore makes no statement that a line of descent does not exist; it states which declared clades cannot share which ancestors.

The verdicts are conditional on the trees as they stand. The ancient samples in Source A are late antique or medieval and lie below the meeting nodes, so none of them bears on the age of a node used here, and no dated ancient genome from the Hijaz or Najd is available. Should one appear and move the age of a node, the test is to be re-applied by the same procedure.

Three further limits attach to the tree itself. A listed SNP may be a false call, or a recurrent or back mutation; one such SNP changes a path count by one, and the shortest path used here has one. The topology may change: new samples can split a branch or re-arrange a polytomy, and a verdict graded here may then move. And the three results are not equally informative. The exclusion at J-FGC12 is the most secure and the least informative, since its meeting is so deep that the result is unsurprising; the informative results are those at J-FGC23310 and J-FGC4302, which are also the ones nearest their thresholds.

STR results come from project members who chose to test and to join; kits were used only where an SNP-confirmed haplogroup placed them on the tree, and rare large allele jumps were controlled by capping, not by modelling. Samples are limited and biased towards those who have tested; for the meeting ages used here that bias does not move the node, and no claim of conservatism is made from it.

The paper refutes and does not confirm. What is not refuted here is not thereby established; it has merely not been reached by this tool. And every verdict is a lower bound: it says below which rung two sides cannot meet, not where they do.

10. Conclusion

This paper has offered a criterion requiring neither classification nor label nor the designation of a reference line: the minimum generation length, computed from the meeting ages of the sides under study and read as exclusion only. Exclusion is monotone, so its verdicts do not depend on which claims are admitted or which subclades represent them; and since every side is taken in turn as holding the measure, and the sides themselves are fixed by a declared rule, none is privileged.

The rule draws sixteen clades from the public record, and by its construction every pair of them is excluded at Hashim by the central estimates of the database that defines them; the findings are the grades. Five units, J-L859 among them, part from all the rest some fourteen centuries before Hashim and are firmly excluded with every other unit at every rung to ʻAdnan. J-FGC7944, J-KM2, J-FGC7 and J-FT38741 part from one another more than six centuries before him and are firmly excluded from one another through Ghalib and probably through Kinana. Inside J-FGC7, J-FGC51966 is firmly excluded through Kaʻb and J-ZS11906 at Hashim; inside J-KM2, J-FTE99194 through Murra. In all, 115 of the 120 pairs are firmly excluded from joint descent from Hashim.

Five pairs are not, all of them between neighbouring units inside one clade. The three units that remain under J-FGC4302 are probably but not firmly excluded from one another at Hashim, and the verdict there rests on the precision assigned to the input ages; one pair inside J-KM2 and one inside J-FGC7944 are undecided. On the wider intervals that Source B publishes for its own nodes, only the exclusions of the deep level reach above Hashim. Whichever side holds the measure, J-L859 is excluded with every other side at every rung from Hashim to ‘Adnan, in both databases, at all five thresholds, under every published mutation rate, on Source B’s own interval, under the systematic shrinkage and the 1950 scenario; an STR clock confirms it through Mudar, and at ‘Adnan when large jumps are capped. These verdicts are lower bounds on where the sides meet, and none of them locates a meeting.

The procedure is general. For any set of competing claims to one eponymous ancestor, a declared rule fixes the claimant clades, the tree fixes their meeting ages, and the ladder of generations turns each age into the rungs at which a pair is excluded. It separates claims in proportion to their distance in the tree, and it is silent on claims that lie close together.

The weakest premise is not the generation length but the mutation rate: a shift of 6–13% in the calibration, within the range of published rates, is enough to move the verdicts on J-FGC23310 and J-FGC4302, and the precision assigned to the ages of the descendant nodes moves them as much. Further tightening therefore requires a calibration of the Y-SNP rate independent of the dating databases, more samples in the sparse sides, and an explicit statement by the databases of their references, rates and calibration methods.

Use of artificial intelligence

An AI assistant was used to parse and sort the data, to write and run the analysis code, and to carry out the statistical computations. The author directed the analysis, reviewed the results and is responsible for the content.

References

Adamov D, Guryanov V, Karzhavin S, Tagankin V, Urasin V (2015) Defining a new rate constant for Y-chromosome SNPs based on full sequencing data. The Russian Journal of Genetic Genealogy 7(1):68–89 (article in English)

Balanovsky O (2017) Toward a consensus on SNP and STR mutation rates on the human Y-chromosome. Hum Genet 136(5):575–590. https://doi.org/10.1007/s00439-017-1805-8

FamilyTreeDNA. Public group projects: member lists of the projects read for the rule of Section 3.6, and the Y-DNA results pages of three projects for STR values (SNP-tested kits only); FamilyTreeDNA Discover. https://discover.familytreedna.com/. Accessed November 2026

Fenner JN (2005) Cross-cultural estimation of the human generation interval for use in genetics-based population divergence studies. Am J Phys Anthropol 128(2):415–423. https://doi.org/10.1002/ajpa.20188

Fu Q, Li H, Moorjani P et al (2014) Genome sequence of a 45,000-year-old modern human from western Siberia. Nature 514:445–449. https://doi.org/10.1038/nature13810

Hammer MF, Behar DM, Karafet TM, Mendez FL, Hallmark B, Erez T, Zhivotovsky LA, Rosset S, Skorecki K (2009) Extended Y chromosome haplotypes resolve multiple and unique lineages of the Jewish priesthood. Hum Genet 126(5):707–717. https://doi.org/10.1007/s00439-009-0727-5

Helgason A, Einarsson AW, Guðmundsdóttir VB et al (2015) The Y-chromosome point mutation rate in humans. Nat Genet 47:453–457. https://doi.org/10.1038/ng.3171

Ibn Hazm, ‘Ali b. Ahmad (d. 456 AH/1064 CE) Jamharat Ansab al-‘Arab

Ibn Hisham, ‘Abd al-Malik (d. 218 AH/833 CE) al-Sira al-Nabawiyya

Ibn Sa‘d, Muhammad (d. 230 AH/845 CE) Kitab al-Tabaqat al-Kabir

Lipson J, Bohrer SG, Cohen-Weinstein S, Cooper L, Revivo D, Staetsky D, Maier P, Runfeldt G, Skorecki K (2025) Revisiting the lineages of the Cohanim using data from next-generation sequencing. bioRxiv preprint. https://doi.org/10.64898/2025.12.08.692646. Accessed October 2026

McDonald I (2021) Improved models of coalescence ages of Y-DNA haplogroups. Genes 12(6):862. https://doi.org/10.3390/genes12060862

Skorecki K, Selig S, Blazer S, Bradman R, Bradman N, Waburton [sic] PJ, Ismajlowicz M, Hammer MF (1997) Y chromosomes of Jewish priests. Nature 385(6611):32. https://doi.org/10.1038/385032a0

Swiel Y, Kelso J, Peyrégne S (2025) Resolving the source of branch length variation in the Y chromosome phylogeny. Genome Biol 26:4. https://doi.org/10.1186/s13059-024-03468-4

TheYtree. Dated J-FGC11 tree. https://www.theytree.com/. The database publishes no version number and no update date. Accessed November 2026

Wang RJ, Al-Saffar SI, Rogers J, Hahn MW (2023) Human generation times across the past 250,000 years. Sci Adv 9(1):eabm7047. https://doi.org/10.1126/sciadv.abm7047

YFull (2026a) YTree v14.07.00 (2 November 2026). Complete J-FGC11 tree. https://www.yfull.com/tree/. Accessed November 2026

YFull (2026b) Age-estimation methodology of the YFull YTree. https://www.yfull.com/faq/what-yfulls-age-estimation-methodology/. Accessed November 2026

YFull (2026c) Definitions of the terms used in the YTree (formed, TMRCA, ybp). yfull.com, FAQ pages. Accessed November 2026

Statements and Declarations

Funding. The author declares that no funds, grants, or other support were received during the preparation of this manuscript.

Competing interests. The author has no relevant financial or non-financial interests to disclose.

Ethics approval. This study analysed only Y-chromosome data already made publicly available by their owners on public genealogy platforms (YFull YTree, TheYtree and FamilyTreeDNA public group projects). No participants were contacted or recruited, no samples were collected, and no individual is named or identified. The author has no institutional affiliation; no ethics committee review was sought.

Consent to participate. Not applicable; no new data were collected from human participants.

Data availability. All data analysed in this study are derived from publicly accessible sources: the YFull YTree (https://www.yfull.com/tree/), FamilyTreeDNA Discover (https://discover.familytreedna.com/) and TheYtree (https://www.theytree.com/), as cited in the text. The compiled tables of clade ages, SNP counts and unit assignments are fully reported in the manuscript. The analysis code and intermediate files are available from the author upon reasonable request.

Code availability. Available from the author upon reasonable request.

Author contributions. S.F.A. conceived the study, compiled the data, directed the analysis and wrote the manuscript.

Computational appendix

The programs behind the computed quantities of this paper are kept together with the data they read: the reader and tree reconstructor for the YFull export, the parser for the TheYtree export, the exhaustive check of meeting ages under every internal reference, the database comparison and stability calibration with its

two tails, the branch SNP-count simulation under each mutation rate, the STR distance and calibration code, and the figure generator. They are available from the author on request.

The dating model. The age of a meeting node X is resolved from well-dated nodes below it. For path i, whose end node has the published age T(i) and the assigned relative error s(i), and which carries n(i) SNPs between its end node and X,

$$A(i) = c \cdot T(i) \cdot (1 + e(i)) + G(i) \cdot r$$

where e(i) is normal with mean 0 and standard deviation s(i); G(i) follows a gamma distribution with shape n(i) + a and unit scale, with a = 0 in the base case; r, the number of years per SNP, is normal with mean 80c and standard deviation 6c; and c = 0.818 ÷ μ rescales Source B's ages and its years per SNP to the mutation rate μ. The paths are combined by inverse-variance weights,

$$T(X) = \Sigma\, w(i) \cdot A(i) \div \Sigma\, w(i), \quad w(i) = 1 \div \mathrm{Var}\, A(i)$$

and the probability at rung k is the share of at least 200,000 draws in which T(X) exceeds T_0 + 35k. The heterogeneity statistic Q is the weighted sum of squared deviations of the path means from their weighted mean. Correlation between paths that share upper branches is ignored. The value of r is taken from Source B itself: the total duration of its branches divided by the total of SNPs listed on them is 81 years per SNP, and the model uses 80.

Table A1 Paths used to resolve the meeting nodes. For each path: the end node, then in brackets its age in Source B, the relative error assigned to that age, and the number of SNPs between it and the meeting node

Meeting node	Paths
J-FGC4302	J-ZS4753 (1400; 10%; 3), J-BY258 (1350; 20%; 6), J-FGC39219 (1500; 30%; 5), J-FGC4320 (1200; 30%; 9), J-FGC19487 (1550; 28%; 7), J-FGC11743 (1350; 24%; 5), J-Y89244 (1100; 24%; 13)
J-FGC4316	J-HU331 (1400; 20%; 3), J-ZS7241 (1550; 20%; 6), J-FGC11722 (1650; 24%; 2), J-FTE36043 (1050; 40%; 1), J-A33246 (1350; 40%; 7), J-ZS11906 (550; 40%; 22), J-BY37632 (1650; 24%; 6)
J-FGC9	J-FGC51966 (1400; 30%; 4), J-Y8962 (1450; 16%; 7), J-FGC13050 (1400; 16%; 23)
J-FGC23310	J-Y8975 (1800; 12%; 5), J-Y10731 (1650; 24%; 15), J-FGC19118 (700; 50%; 23), J-FGC10579 (1650; 16%; 4)
J-FGC10562	J-FT195589 (800; 30%; 9), J-L861 (1300; 20%; 4), J-Y207222 (750; 40%; 12), J-FT352159 (800; 50%; 7), J-MF19113 (850; 40%; 3), J-FT159397 (1650; 24%; 7)
J-FGC5	J-FT198133 (1200; 40%; 6), J-ZS10807 (1600; 24%; 4), J-FGC4285 (1400; 24%; 1), J-Y5320 (1600; 10%; 2)
J-KM2	J-Y5323 (1800; 10%; 2), J-Y485299 (1800; 16%; 1), J-BY2 (1600; 20%; 2), J-FT198133 (1200; 40%; 10), J-ZS10807 (1600; 30%; 8), J-FGC4285 (1400; 24%; 5), J-Y5320 (1600; 10%; 6), J-FGC23914 (1200; 16%; 7), J-Y19266 (1650; 30%; 10), J-Y60236 (1150; 30%; 22), J-Y8968 (1700; 20%; 5), J-FGC50967 (1550; 24%; 5), J-Y20496 (1450; 24%; 6)

J-FGC4316 uses in addition the seven paths of J-FGC4302, each with two SNPs more. J-FGC9 uses in addition the fourteen paths of J-FGC4316, each with one SNP more. J-KM2: its three immediate branches and, below J-FGC5 and J-FGC1713, the first nodes younger than their own formation that hold eight samples or more.

Table A2 Sensitivity of the grades to the model: the highest rung at which the exclusion is firm / probable

Meeting node	As reported	a = 0.5	a = 1	Errors halved	Errors × 1.5	Unweighted mean	One path removed: firm through
J-FGC4302	— / Hashim	— / Hashim	Hashim / Qusayy	Hashim / Hashim	— / —	— / Qusayy	—
J-FGC4316	Hashim / Murra	Qusayy / Ka‘b	Qusayy / Lu’ayy	Hashim / Qusayy	Hashim / Qusayy	Qusayy / Ka‘b	Hashim to Qusayy
J-FGC9	Ka‘b / Ghalib	Lu’ayy / Fihr	Ghalib / Fihr	Murra / Lu’ayy	Murra / Ghalib	Lu’ayy / Fihr	Qusayy to Lu’ayy
J-FGC23310	Ghalib / Kinana	Fihr / Kinana	Fihr / Mudrika	Fihr / Kinana	Ka‘b / al-Nadr	al-Nadr / Ilyas	Murra to Fihr

The effects are not all in one direction. A positive offset and an unweighted mean give exclusions at least as strong at every node. Narrower errors strengthen the exclusion at J-FGC4302 and J-FGC23310 and weaken it at J-FGC4316 and J-FGC9, because they shift the weight among the paths. Errors half as wide again weaken it at every node, and at J-FGC4302 they remove the probable grade at Hashim.

Table A3 The intervals Source B publishes for the meeting nodes (YFull v14.07.00) beside the resolved ages at Source B’s own rate, with the highest rung whose threshold lies below the lower end of each

Node (YFull label)	YFull age	YFull 95% interval	Model median	Model 95% interval	Reach on YFull interval	Reach on model interval
J-FGC4302 (J-Y10966)	1700	1500–2000	1802	1557–2072	none	none
J-FGC4316 (J-Y8975)	1800	1600–2000	1908	1697–2132	Hashim	Qusayy
J-FGC9 (J-Y9274)	2200	1600–2800	2033	1828–2250	Hashim	Ghalib
J-FGC23310	2200	1600–2800	2238	1888–2619	Hashim	Fihr
J-FGC12 (J-ZS2121)	3000	2800–3300	—	—	‘Adnan	—

Reach: the highest rung whose threshold at g = 35 lies below the lower end of the 95% interval. Both columns are for Source B’s rate alone; the grades of the paper require all three published rates.

Table A4 The three small levels inside J-FGC7944 and J-KM2: probability that the divergence predates each of the three nearest ancestors (g = 35), under three published mutation rates

Level	Pair	Ancestor	Fu 0.76	YFull 0.818	Helgason 0.871	Grade
J-FGC10562	FT195589 with Y207222	Hashim	0.75	0.49	0.25	undecided
J-FGC10562	FT195589 with Y207222	Qusayy	0.60	0.32	0.13	undecided
J-FGC10562	FT195589 with Y207222	Murra	0.44	0.18	0.06	undecided
J-FGC5	Y8537 with FT198133	Hashim	0.97	0.86	0.64	undecided
J-FGC5	Y8537 with FT198133	Qusayy	0.92	0.72	0.45	undecided
J-FGC5	Y8537 with FT198133	Murra	0.82	0.54	0.27	undecided
J-KM2	FTE99194 with Y8537 and FT198133	Hashim	1.00	1.00	1.00	firm
J-KM2	FTE99194 with Y8537 and FT198133	Qusayy	1.00	1.00	0.99	firm
J-KM2	FTE99194 with Y8537 and FT198133	Murra	1.00	1.00	0.96	firm
J-KM2	FTE99194 with Y8537 and FT198133	Ka‘b	1.00	0.99	0.91	probable
J-KM2	FTE99194 with Y8537 and FT198133	Lu’ayy	1.00	0.98	0.83	probable
J-KM2	FTE99194 with Y8537 and FT198133	Ghalib	1.00	0.96	0.72	undecided

Six, four and thirteen paths (Table A1). Above the rungs shown every probability is lower still.

Table A5 Check of the reconstructed trees: meeting ages under every choice of reference. Every dated subclade of one side was paired with every dated subclade of the other. Subclades per side: Source A — L859 44, FGC7944 40, ZS4753 50, FGC4309 32; Source B — 95, 160, 70, 68

Pair of sides	Source A: combinations	Source A: distinct meeting ages	Source B: combinations	Source B: distinct meeting ages
L859 – FGC7944	1,760	one (3100)	15,200	one (3000)
L859 – ZS4753	2,200	one (3100)	6,650	one (3000)
L859 – FGC4309	1,408	one (3100)	6,460	one (3000)
FGC7944 – ZS4753	2,000	one (2230)	11,200	one (2200)
FGC7944 – FGC4309	1,280	one (2230)	10,880	one (2200)
ZS4753 – FGC4309	1,600	one (1790)	4,760	one (1700)

Table A6 Stability of each node verdict: shrinkage that would overturn the exclusion / frequency of a disagreement that large in the overturning direction, all nodes (deep nodes) / band

Meeting node	Rung	Source A	Source B
J-FGC12	‘Adnan	29% / 13% (10%) / moderate	27% / 11% (0%) / high
J-FGC12	Mudar	33% / 11% (7%) / high	30% / 9% (0%) / high
J-FGC12	Ilyas	34% / 10% (6%) / high	32% / 8% (0%) / high
J-FGC12	Mudrika	35% / 10% (6%) / high	33% / 8% (0%) / high
J-FGC12	Kinana	37% / 9% (3%) / high	35% / 7% (0%) / high
J-FGC12	Fihr	41% / 7% (1%) / high	39% / 6% (0%) / high
J-FGC12	Murra	45% / 5% (0.6%) / high	43% / 5% (0%) / high
J-FGC12	Hashim	50% / 4% (0%) / high	48% / 3% (0%) / high
J-FGC23310	‘Adnan	2% / 46% (74%) / low	0.4% / 33% (17%) / low
J-FGC23310	Mudar	6% / 40% (62%) / low	5% / 29% (8%) / low
J-FGC23310	Ilyas	8% / 37% (54%) / low	7% / 27% (6%) / low
J-FGC23310	Mudrika	10% / 34% (49%) / low	8% / 24.7% (4%) / moderate
J-FGC23310	Kinana	13% / 30% (41%) / low	12% / 22% (3%) / moderate
J-FGC23310	Fihr	17% / 23% (24.8%) / moderate	16% / 19% (0.6%) / moderate
J-FGC23310	Murra	24% / 16% (15%) / moderate	23% / 13% (0%) / moderate
J-FGC23310	Hashim	30% / 12.05% (8%) / moderate	29% / 10% (0%) / high
J-FGC4302	Murra	5% / 43% (64%) / low	not excluded
J-FGC4302	Qusayy	9% / 35% (52%) / low	4% / 30% (8%) / low
J-FGC4302	Hashim	13% / 30% (41%) / low	8% / 25.3% (4%) / low

Table A7 The fastest Y-SNP mutation rate ($\times 10^{-9}$ per site per year) at which the exclusion remains firm, by meeting node and rung, for g = 35 / g = 30. Published rates: 0.76, 0.818 and 0.871

Ancestor	k	J-FGC4302	J-FGC4316	J-FGC9	J-FGC23310	J-FGC12 †
Hashim	3	0.83 / 0.84	0.90 / 0.91	0.97 / 0.98	1.01 / 1.02	1.38 / 1.39
Qusayy	5	0.79 / 0.81	0.86 / 0.88	0.93 / 0.94	0.97 / 0.98	1.32 / 1.34
Murra	7	0.76 / 0.78	0.83 / 0.84	0.89 / 0.91	0.93 / 0.95	1.27 / 1.29
Ka‘b	8	0.75 / 0.76	0.81 / 0.83	0.87 / 0.89	0.91 / 0.93	1.24 / 1.27
Lu’ayy	9	0.73 / 0.75	0.79 / 0.81	0.85 / 0.88	0.89 / 0.91	1.22 / 1.25
Ghalib	10	0.72 / 0.74	0.78 / 0.80	0.84 / 0.86	0.87 / 0.90	1.19 / 1.23
Fihr	11	0.70 / 0.73	0.76 / 0.79	0.82 / 0.85	0.86 / 0.88	1.17 / 1.21
al-Nadr	13	< 0.70 / 0.70	0.74 / 0.76	0.79 / 0.82	0.83 / 0.86	1.13 / 1.17
Kinana	14	< 0.70 / < 0.70	0.72 / 0.75	0.78 / 0.81	0.81 / 0.84	1.11 / 1.15
Mudrika	16	< 0.70 / < 0.70	0.70 / 0.73	0.75 / 0.78	0.78 / 0.82	1.07 / 1.11
Ilyas	17	< 0.70 / < 0.70	< 0.70 / 0.71	0.74 / 0.77	0.77 / 0.80	1.05 / 1.10
Mudar	18	< 0.70 / < 0.70	< 0.70 / 0.70	0.72 / 0.76	0.76 / 0.79	1.03 / 1.08
‘Adnan	21	< 0.70 / < 0.70	< 0.70 / < 0.70	< 0.70 / 0.72	0.72 / 0.76	0.98 / 1.03

† Not resolved from paths: computed from the lowest central meeting age of the deep level (2900 years), reduced by the 9.1% of Section 7.3. “< 0.70”: not firm at any rate in the range examined.