

Materials and Methods

Charles Kemp & Terry Regier

1	Cross-cultural data	1
2	Cousins analyses	3
3	Representation language	5
4	Generating kinship categories: analyses from main text	7
5	Generating kinship categories: cousins analyses	10
6	Generating kin classification systems	12
7	Complexity of kin classification systems	14
7.1	Finding the simplest representation of a kin classification system	14
7.2	Alternative complexity measures	16
8	Need probabilities	18
9	Adjusting for a stable population	18
10	Analyses from main text	20
10.1	Figures 4A and 4B: partitions of the family tree	20
10.2	Figure 4C: grandparents, grandchildren, siblings, aunts, uncles and niblings . .	21
10.3	Figures 5B and 5C: permutation analysis	23
10.4	Figure 5D: conjunctive vs disjunctive categories	24

1 Cross-cultural data

We used an electronic version of Murdock’s 1970 data set (60). The original data file contains tables for eight kinds of kin types: grandparents, grandchildren, uncles, aunts, niblings

(nephews and nieces), siblings, cousins, and siblings-in-law.¹ Working with a family tree that included all of these kin types proved to be computationally intractable, and the analyses in the main text focus on kin types that could be reached by traversing up to three parent/child links in the family tree. This criterion includes six of Murdock's tables but leaves out his data for cousins and siblings-in-law. Analyses that include cousins are described in the following section.

Although Murdock presented separate data tables for different kinds of kin types, we needed to merge these tables to produce a single kin classification system for each language. In some cases Murdock's variable descriptions indicate that a kin type included in one table is described using the same term as a kin type from elsewhere in the family tree. For example, the table for grandchildren includes cases where the same term is used to refer to grandparents and grandchildren. We allowed for all of these cases, but assumed that in all other cases the terms for kin types in different tables were distinct. This assumption is not strictly true—for example, kin classification systems of the Crow type (2) often use the same term for grandmothers and aunts, but this overlap is not noted in Murdock's data.

A second limitation of Murdock's data is that it does not specify kin terms for parents and children. We therefore assumed that every language has four categories which can be glossed as "mother," "father," "daughter" and "son." Murdock's codes for aunts and uncles do specify cases where one or more aunts are grouped with mother and where one or more uncles are grouped with father, and we expanded the "mother" and "father" categories as needed in these cases. Allowing two categories for "mother" and "father" is consistent with the claim that distinguishing mothers and fathers is a cultural universal (3, 28), although this claim has been disputed (61, 62). Some languages, however, do not include terms for "son" and "daughter" (24),

¹A small number of entries in the aunt table appeared to be incorrect—they are inconsistent with the variable descriptions that accompany the electronic data set and they do not match the information in Murdock's original article (27). We removed these errors and reconciled the electronic data set with the original article by changing all occurrences of (7, 10) and (8, ·) to (7, 9) and (8, 10) respectively.

and the kin classification systems used for our analyses are therefore inaccurate in some cases.

Although Murdock’s data set is not as complete as we would have liked, we believe that our method for converting his tables into complete kin classification systems is theoretically neutral with respect to our conclusions. Note also that the subtree analyses in Figure 4 are unaffected by any possible distortions that resulted from merging Murdock’s tables or adding assumptions about parent and children categories. Each subtree analysis is based on the data contained in just one of Murdock’s tables, and in all cases the results provide evidence that attested kin classification systems tend to be near-optimal.

Murdock’s data set includes 566 languages in total, and one or more of the codes for the kin types we considered were missing for 60 of these languages. An additional 19 languages had codes that could not be converted into a partition. For example, the grandchild codes for some languages indicate that “grandparental terms are applied to grandchildren but not always self-reciprocally,” and this description does not fully specify how Alice and Bob in Figure 2 refer to their grandchildren. After removing all languages with codes that were missing or underspecified, we were left with complete kin classification systems for 487 languages.

2 Cousins analyses

To explore whether our theory extends to categories for cousins, we worked with the family trees shown in Figure S1. Many discussions of kin classification refer to six different families of systems known as Sudanese, Hawaiian, Eskimo, Iroquois, Omaha and Crow. Murdock’s data for cousins assign each attested system to one of these families. Figure S2 shows representatives of each family and demonstrates that classification systems defined over the trees in Figure S1 are extensive enough to capture some important differences between these families. For the six systems shown, male and female speakers use the same kinship categories for the

kin types in Figure S1, and the trees in Figure S2 are therefore based on a speaker labeled Ego of indeterminate sex.

Although the tree we consider is large enough to capture some important characteristics of the six families of systems, our analyses do not address all aspects of cousin classification. First, the tree does not include descending generations such as daughters, sons, nieces and nephews. Omaha and Crow systems include terms that group some of these kin types with kin types such as parents, aunts, and uncles, and these terms are typically used differently by male and female speakers. We omitted descending generations from the trees in Figure S1 because Murdock's data do not specify kinship terms for daughters and sons, or for nieces and nephews of a female speaker. Second, the Omaha and Crow systems in Figure S2 only hint at the fact that these systems tend to have categories that include unlimited numbers of kin types. For example, Omaha-type systems tend to have a category that groups MBD with M, MZ, MFFM and other female members of one's mother's patrilineage, and an infinite tree would be needed to capture the extension of this category. As discussed in the following section, our representation language can express categories like "female member of my mother's patrilineage," but an analysis based on the finite trees in Figure S1 does not constitute a strong test of this capacity. Note, however, that the simple trees in Figure S1 are sufficient to capture some important properties of Omaha and Crow systems, including the basic fact that both of these systems group cousins (e.g. MBD) with kin types from different generations.

For present purposes we will refer to a partition of the family trees in Figure S1 as a *cousin system*. Following the approach described in the main text, we carried out one analysis that compares attested cousin systems to a large space of alternative systems, and a second fine-grained analysis that compares attested cousin systems to permutations of these systems. The results of the first analysis are shown in Figure S3. Black circles represent attested systems, and the light grey circles represent 100 million hypothetical systems that were sampled from a

space of systems described in Section 5 below. As for the plot in Figure 4A, attested systems tend to appear near the optimal frontier.

The permutations considered for the permutation analysis were constructed by exchanging the five chunks shown in Figure S4A. The results are shown in Figure S4B, which indicates that attested systems tend to score better than permutations of these systems. To avoid creating unnatural categories, we excluded permutations of an attested system that move only part of a category that extends beyond the five chunks. For example, Omaha-type systems include a category that groups MBD with MZ, and that therefore extends beyond the five chunks in Figure S4A. We therefore ruled out permutations of Omaha-type systems that exchange MBD and MBS with some other chunk. For each attested system, we also ruled out permutations which left that system unchanged. For example, Eskimo-type systems group the eight cousins in Figure S4A into a single category, and we ruled out permutations of Eskimo-type systems (e.g. exchanging MBD and MBS with MZD and MZS) that simply rearrange the members of this category.

The results in Figures S3 and S4 parallel the results in Figures 4 and 5 and suggest that the conclusions described in the main text extend to kin classification systems which include cousins. In particular, Figure S3 suggests that attested cousin systems are near-optimal with respect to the space of all possible systems, and Figure S4 suggests that attested cousin systems tend to lie closer to the optimal frontier than do simple transformations of these systems.

3 Representation language

Our representation language corresponds to a working hypothesis about the conceptual resources that contribute to mental representations of kinship categories. The rules in Figure 3B are formulated in a representation language that includes conjunction (\wedge) and disjunction (\vee),

and that specifies relative products using an existential quantifier (\exists). As illustrated in Figure 2C, relative products are useful for capturing categories like $\text{sister}(\cdot, \cdot)$:

$$\text{sister}(x, y) \leftrightarrow \exists z \text{ daughter}(x, z) \wedge \text{PARENT}(z, y)$$

where this definition indicates that x is the sister of y if and only if there is some person z such that x is the daughter of z and z is the parent of y . Note that this definition allows a woman x to qualify as her own sister. Because kinship categories are used to refer to individuals other than one's self, we assume that identity relationships are ruled out.

The primitives in Figure 3A and the composition rules in Figure 3B are relatively standard and are consistent with previous formal accounts of kinship semantics (30, 63). Other representational choices are possible, however, and given any candidate representation language, our general approach can be used to compute which kin classification systems are near-optimal with respect to that language.

An adequate representation language should be able to capture kinship terms for both near and distant relatives. Because Murdock's data specify kin terms for near relatives, all of our analyses focus on near rather than distant kin. Our representation language, however, can be used to define kin categories that include distant relatives, including kin categories that include an unbounded number of individuals. Consider, for example, the category that includes "women in my mother's patrilineage." This category is characteristic of Omaha-type systems (2) and can be defined in our language as follows:

$$\begin{aligned} \text{patriline}(x, y) &\leftrightarrow \text{father}^+(x, y) \\ \text{patrilineal_descendant}(x, y) &\leftrightarrow \text{patriline}(y, x) \\ \text{mothers_patriline}(x, y) &\leftrightarrow \exists z \text{ patriline}(x, z) \wedge \text{mother}(z, y) \\ \text{mothers_patrilineage}(x, y) &\leftrightarrow \exists z \text{ patrilineal_descendant}(x, z) \wedge \text{mothers_patriline}(z, y) \end{aligned}$$

$woman_in_mothers_patrilineage(x, y) \leftrightarrow mothers_patrilineage(x, y) \wedge FEMALE(x)$

As this example suggests, the transitive closure operator can be used to define concepts that include an unlimited number of individuals. For example, the concept $patriline(x, y)$ defined above includes y 's father, y 's father's father, y 's father's father's father, and so on. Although the transitive closure operator allows concise definitions of some terms for distant kin, the simple representation language used in this paper may need to be refined or extended in order to represent all attested terms for distant kin. Our general approach, however, can still be applied even if the specific language used here is adjusted.

Although our representation language relies on genealogical primitives such as $PARENT(\cdot, \cdot)$, one important approach to kinship does not rely on genealogical primitives of this kind. Instead, certain kin terms are treated as primitives, and other kin terms are expressed as combinations of these primitive terms. Read (2007) argues for an approach of this kind (49) and suggests that it can handle cases where kin terms are used to refer to individuals who bear no genealogical relationship to the speaker. In keeping with much of the literature on kinship, we have focused on genealogical primitives, but the formal apparatus we developed could be applied to representation languages that do not take genealogical primitives as a starting point.

4 Generating kinship categories: analyses from main text

Our analyses make use of a large collection \mathcal{C} of possible kinship categories, including almost all categories found in the Murdock data. It will be convenient to distinguish among three aspects of the meaning of a category. Consider the category corresponding to the English term $mother(\cdot, \cdot)$. The *extension* of the category is the set of pairs (x, y) where x and y both belong to

the set of 58 individuals in Figure 2 (56 relatives plus Alice and Bob), and x is the mother of y . For example, if Abigail is the mother of Ada and Ada is the mother of Alice then the extension of $\text{mother}(\cdot, \cdot)$ includes the pairs (Abigail, Ada) and (Ada, Alice). A category *intension* is a definition constructed using the resources in Figure 3. For example, in Figure 2C the intension of $\text{sister}(\cdot, \cdot)$ is

$$\text{sister}(x, y) \leftrightarrow \exists z \text{ daughter}(x, z) \wedge \text{PARENT}(z, y)$$

but an alternative intension is

$$\text{sister}(x, y) \leftrightarrow \text{sibling}(x, y) \wedge \text{FEMALE}(x).$$

Finally, the *ego-relative extension* of $\text{mother}(\cdot, \cdot)$ is the set of all relatives x in the trees of Figure 2 such that $\text{mother}(x, \text{Alice})$ or $\text{mother}(x, \text{Bob})$. Note that this set is different from the set of mothers: Abigail is a mother, but does not appear in the ego-relative extension of $\text{mother}(\cdot, \cdot)$ because she is not a mother of Alice or Bob. Ego-relative extensions will be important for our purposes because each partition in the Murdock data corresponds to a collection of these extensions.

Because there are 58 individuals in Figure 2, there are 3364 possible pairs, and the number of possible category extensions is 2^{3364} . There are infinitely many possible intensions: note, for example, that $\text{mother}(\cdot, \cdot)$ could be defined as

$$\text{m1}(x, y) \leftrightarrow \text{PARENT}(x, y) \wedge \text{FEMALE}(x)$$

$$\text{m2}(x, y) \leftrightarrow \text{m1}(x, y) \wedge \text{FEMALE}(x)$$

$$\text{mother}(x, y) \leftrightarrow \text{m2}(x, y) \wedge \text{FEMALE}(x)$$

and the same approach could be used to define $\text{mother}(\cdot, \cdot)$ using any given number of rules. Finally, the number of distinct ego-relative extensions is 2^{56} , because the trees in Figure 2 include 56 relatives.

We worked with a finite set of extensions, intensions, and ego-relative extensions generated by a depth-three expansion of the primitives in Figure 3A using the rules in Figure 3B. A depth-three expansion generates virtually all of the attested categories in the Murdock data, but a depth-two expansion does not adequately cover the space of attested categories. The depth-three expansion was carried out as follows. We began with the eight primitive categories—the depth 0 set—and created a depth 1 set which includes all intensions that can be constructed by applying a single rule in Figure 3. For example, the depth 1 intensions include the conjunction $\text{mother}(x, y) \leftrightarrow \text{PARENT}(x, y) \wedge \text{FEMALE}(x)$. After creating the depth 1 set, we pruned all intensions that produced the same extension as some other intension but had greater *weight*. The weight of an intension is recursively defined as one plus the weight of each of its components, where primitive categories have weight zero. For example, the intension $\text{C}(x, y) \leftrightarrow \text{PARENT}(x, y) \wedge \text{PARENT}(x, y)$ (weight 1) was pruned because $\text{PARENT}(x, y)$ generates the same extension but has weight 0. We then created intensions at depths 2 and 3, in each case combining intensions from all previous depths then pruning intensions as already described. For the depth 3 expansion only we carried out an additional stage of pruning, and removed all intensions that generated the same ego-relative extension as another intension but had greater weight. Finally, we identified all pairs of extensions which were mutual inverses (e.g. the pair of extensions corresponding to $\text{PARENT}(\cdot, \cdot)$ and $\text{CHILD}(\cdot, \cdot)$), and added intensions indicating that both members of the pair could be defined as inverses of the other. At the end of this process we were left with a collection \mathcal{C} of around 116,000 intensions which produced around 97,000 distinct extensions, which in turn produced around 71,000 distinct ego-relative extensions.

The Murdock dataset contains complete kin classification systems for 487 languages and these systems include 173 distinct categories. The ego-relative extensions generated by the depth 3 expansion described include 164 of these categories. The remaining nine categories are rare: seven of them appear only once in the Murdock data, and two of them appear five times each. These nine rare ego-relative extensions could be generated by expanding to a depth greater than three—note that there is some finite depth after which all possible ego-relative extensions will have been generated. To keep our analysis tractable we terminated the expansion at depth three, and removed all languages from the Murdock set that use one or more of the nine rare categories. There are nine such languages and after they are removed we are left with 478 languages that represent 401 distinct types of kin classification systems.

Most categories in the Murdock data set have conjunctive definitions—if disjunctions are ruled out, then 159 of the 173 categories in the Murdock data set are generated at depth 3 or earlier in the expansion process already described, and each of the excluded categories appears at most 5 times in the full Murdock data set. It is not surprising that combinations of conjunctions and disjunctions are sufficient to capture all of the categories in the Murdock data, but it is notable that a depth-three expansion can account for almost all of the categories. This result suggests that the representational machinery in Figure 3 may allow a concise formulation of nearly all attested kinship categories.

5 Generating kinship categories: cousins analyses

As just described, the collection \mathcal{C} used for all analyses in the main text was constructed using a depth-three expansion of the primitives in Figure 3A. Some attested cousin systems include categories that can only be defined using a depth-five expansion of the primitives in Figure 3A, and carrying out a full depth-five expansion proved to be computationally intractable. Our

cousins analyses therefore make use of a more restricted collection of categories \mathcal{C} that was designed to include all simple categories that might play a role in analyzing systems of aunts, uncles, and cousins. To enumerate this collection, our first step was to create 5 groups that include 11 categories each:

- 1) Group P: parent, mother, father, same sex parent, different sex parent, parent of female, parent of male, mother of female, mother of male, father of female, father of male
- 2) Group C: child, daughter, son, same sex child, different sex child, child of female, child of male, daughter of female, daughter of male, son of female, son of male
- 3) Group Sl: sibling, sister, brother, same sex sibling, different sex sibling, sibling of female, sibling of male, sister of female, sister of male, brother of female, brother of male
- 4) Group An: ancestor, . . .
- 5) Group De: descendant, . . .

The first three groups were created by taking a relation R and constructing the following eleven categories: R , female R , male R , same sex R , different sex R , R of female, R of male, female R of female, female R of male, male R of female, and male R of male. The final two groups were constructed by taking the transitive closure of all relations in groups 1 and 2.

Next we considered all natural ways to combine the five groups to define categories that include kin types in Figure S1. The approach taken is summarized by Figure S5. Figure S5A shows combinations involving the parent, child, and sibling groups. P of P indicates that two categories from group P are combined using a relative product: for example, “parent of parent,” “father of parent,” and “father of same sex parent” can all be generated in this way. C of P of P indicates that one of the resulting categories is then combined with a category from group C : for example, “child of father of parent” is one category that might result. Figure S5B shows

combinations based on groups An and De. Collection \mathcal{C} includes the 5 basic groups along with all categories that can be generated by one of the combinations in Figure S5A or S5B. To complete collection \mathcal{C} , we added categories produced by all two-element conjunctions that combine $\text{FEMALE}(\cdot)$, $\text{MALE}(\cdot)$, $\text{SAMESEX}(\cdot, \cdot)$, or $\text{DIFFSEX}(\cdot, \cdot)$ with one of the categories already belonging to set \mathcal{C} .

The final version of collection \mathcal{C} includes intensions that generate 747 different categories defined over the 32 individuals in Figure S1. These 747 categories include 67 of the 74 categories that appear in the Murdock data. The most common attested category that does not appear in \mathcal{C} includes Alice’s mother and mother’s brother’s daughter along with the same two relatives with respect to Bob. The six remaining attested categories that do not appear in \mathcal{C} are relatively rare, and each one appears in fewer than 10 attested systems. In total, 38 of the 506 attested cousin systems include one or more categories that do not belong to collection \mathcal{C} . Because we cannot compute the complexity of these systems, we removed them from the Murdock data set before carrying out the analyses that generated Figures S3 and S4B. The 468 attested systems that remain include examples of all six families shown in Figure S2.

Although the analyses in the main text and our cousins analyses used different sets of categories \mathcal{C} , the algorithms used for these analyses were otherwise identical. For simplicity, the following sections explain these algorithms by referring to the analyses in the main text, but the same ideas apply to the cousins analyses unless mentioned otherwise.

6 Generating kin classification systems

Having generated a set \mathcal{C} of ego-relative category extensions, we now want to enumerate all possible kin classification systems (i.e. all partitions of the 56 relatives in Figure 2) that can be constructed from these ego-relative category extensions. This challenge can be formulated

as the task of finding all solutions to an *exact cover* problem. Given a set of elements E (e.g. the set of 56 relatives) and a collection of subsets S (e.g. the collection of 71,000 ego-relative extensions described in a Section 4) the exact cover problem is to find a subcollection S^* of S such that each item in E belongs to exactly one subset in S . The Dancing Links algorithm (64) is an efficient method for enumerating solutions to exact cover problems, and we used this algorithm to generate the hypothetical kin classification systems plotted in Figure 4B. The grey points in Figure 4B include all 831,732,650 possible kin classification systems that can be constructed using the 121 kinship categories that appear more than twice in the Murdock data (Figure 4B).

Figure 4A, however, is based on the space of all kin classification systems that can be constructed from the categories in collection \mathcal{C} , and this space is too large to enumerate. We therefore generated a large sample of systems from this space. To sample a system we sampled a kinship category at random, then sampled another category at random from all candidates that did not overlap with the categories already sampled, and continued in this way until we had either a complete partition or a partition that could not be completed using the 71,000 categories in our data set. Incomplete partitions were discarded and we repeated the process until a complete partition was generated.

Because we focus on kin classification systems that correspond to partitions of the 56 relatives in Figure 2, the number of possible systems is the 56th Bell number, or 6.8×10^{54} . Although this collection of systems is vast, systems outside the collection are logically possible, including systems with no terms or multiple terms for some relatives. Systems with no terms for some relatives are penalized by our theory, because the cost of referring to these relatives is effectively infinite. Systems with multiple terms for some relatives will also tend to be penalized, because the communicative cost measure favors systems with small categories, and allowing categories to overlap will tend to increase the average size of the categories in a sys-

tem. Although there are some languages with kin classification systems that do not correspond to partitions (12), all of the kin classification systems in the Murdock data set are partitions, and we therefore focused on a space of possible systems that includes only partitions.

7 Complexity of kin classification systems

Having generated a set of kin classification systems, we now want to find a minimum-complexity representation for each of these systems. We took the complexity of a kin classification system to be the smallest number of rules needed to define each category in the system. We first describe how we found the minimal representation for a kin classification system with respect to this complexity measure, and then discuss other complexity measures that could be considered.

7.1 Finding the simplest representation of a kin classification system

A kin classification system is represented as a collection of categories, and each category is associated with one or more intensions. Some of these intensions may refer to categories that do not appear overtly in the kin classification system: for example, the definitions in Figure 2D refer to the category $sister(\cdot, \cdot)$ which is latent with respect to the Northern Paiute partition. Our goal is to choose a set of intensions where each category included is ultimately grounded in terms of the eight primitives, and where the total complexity of the set is minimized.

We use a depth-first strategy to search for the best representation for a given system that can be constructed using the intensions from the collection \mathcal{C} described in Section 4. At any stage in the search, the “active set” includes all categories that still stand in need of a definition. The search begins with an active set that includes all categories in the partition under consideration. We then remove a category from the active set, choose one definition for this category from \mathcal{C} , and supplement the active set with categories that appear in this definition but have not previ-

ously been added to the active set.² We continue in this fashion, and once the active set becomes empty we have generated a candidate system of rules. If this system is of lower complexity than the current champion, we store it, then backtrack and consider alternative definitions for the categories most recently removed from the active set. The search terminates once we have generated n candidate rule systems or once the set of possible systems is exhausted. For the most computationally demanding analyses (Figures 4A, 4B, S3 and S7) we set $n = 10,000$, and for all remaining analyses we set $n = 100,000,000$. In many cases the search exhausts all possible candidate sets before reaching the threshold set by n , and we know that we have identified the best possible system of definitions given the intensions in collection \mathcal{C} . In other cases, however, the threshold is reached, and we end up with a system of definitions that may not be minimal.

There is a second reason why the complexity scores we compute sometimes overestimate the true scores. When generating intensions, extensions, and ego-relative extensions, recall that we removed intensions which generated the same extension as an alternative with lower weight. For example, $C(x, y) \leftrightarrow \text{PARENT}(x, y) \wedge \text{PARENT}(x)$ would be removed because it generates the same extension as $\text{PARENT}(x, y)$ but has greater weight. This strategy is a useful heuristic because intensions with large weight will tend to substantially increase the complexity of any partition to which they belong. Note, however, that intensions with relatively large weight may sometimes turn out to be cheap when considered in the context of an entire partition. Consider, for example, the English partition in 2A. The weight of

$$\text{aunt}(x, y) \leftrightarrow \exists z \text{ sister}(x, z) \wedge \text{PARENT}(z, y)$$

incorporates the weight of the category $\text{sister}(\cdot, \cdot)$, but this category effectively comes for free

²Some simple strategies lead to improved results in practice. When removing a category from the active set, we always choose the extension with the smallest possible number of definitions. When choosing a definition for the current category, we prefer definitions that use categories which have previously been added to the active set and therefore “come for free.”

because it must be defined anyway as part of the English kin classification system. Although our algorithm successfully computes the minimal complexity score for English, there may be other cases where our complexity scores are overestimates because our set of intensions is incomplete.

7.2 Alternative complexity measures

Our complexity measure is simple and intended to be fairly intuitive, but alternatives could be considered. One alternative is to measure the complexity of a set of rules by counting the total number of *symbols* that it contains rather than the total number of rules. Because each rule in Figure 3B has at most two concepts on the right hand side, the total number of symbols required to express a set of rules will be highly correlated with the complexity measure adopted here.

Another possible approach is to abandon a representation language entirely and to define the complexity of a kin classification system as the number of terms (i.e. named kinship categories) that it contains. We will refer to this approach as the *direct* measure of complexity. The direct measure has the virtue of simplicity but suffers from at least two limitations. First, the representation language used in this paper provides a way to formulate the semantics of kinship terms. If this representation language is abandoned some alternative account of kinship semantics must be provided. Some such account is needed, for example, to explain why many logically possible kinship categories are unattested. Our complexity measure predicts that categories with lengthy definitions in our representation language are dispreferred, but the direct measure suggests that any conceivable category should be viable as long as it is informative. Second, even if some alternative account of individual kinship categories can be supplied, the direct measure is insensitive to whether systems of these categories cohere well as a whole. In contrast, our complexity measure suggests that certain categories are especially likely to appear together. For example, any system that includes the category “woman’s daughter’s child” is

relatively likely to include the category “maternal grandmother,” because the first category can be concisely defined as the reciprocal of the second.

We wished to determine whether the Murdock data exhibit system-internal coherence of the sort that would be captured by our measure but not the direct measure. If the data do exhibit such system-internal coherence, then this coherence should be disrupted by swapping categories between systems, and this disruption should result in increased complexity. To test this prediction we took the original Murdock data set and created randomized data sets by shuffling categories back and forth between systems. To create these randomized data sets we identified cases where the union of two categories matched another category (e.g. the union of `younger_sister` and `older_sister` is `sister`), and where the union of two categories matched the union of two different categories (e.g. the union of `younger_sister` and `older_sister` matches the union of `sister_of_male` and `sister_of_female`). Cases like these allow swaps between classification systems: if one partition includes `younger_sister` and `older_sister` and another includes `sister`, then these categories can be swapped while ensuring that the categories in both systems are still complete and non-overlapping. Each randomized data set was created by carrying out 100,000 randomly chosen swaps. Note that the randomized data sets include exactly the same categories with exactly the same frequencies as the Murdock data—the only difference is that these categories are organized into different systems.

Figure S7 confirms the prediction that the shuffled hypothetical systems tend to be more complex with respect to our complexity measure than the attested systems from which they were derived. The shuffling analysis therefore suggests that the Murdock data do exhibit system-internal coherence, and that our complexity measure is sensitive to this coherence. Note that the shuffling analysis controls for differences in communicative cost and in complexity as measured by the direct approach. We can therefore conclude that our measure captures aspects of conceptual complexity present in the Murdock data that go beyond these factors.

8 Need probabilities

The probabilities in Figure 3C are based on corpus statistics for English and German. Figure S6 shows that the results for the two languages are qualitatively similar, and we therefore averaged these results to create the probabilities shown in Figure 3C.

The results for English are based on the relative frequencies of kin expressions of the form “my grandmother,” “my aunt,” “my mother,” “my sister,” “my daughter,” “my granddaughter” and the like. The statistics also include counts for less formal variants of these phrases—for example, the count for mother includes counts for “my mom,” “my mommy” and “my mama” in addition to “my mother.” The statistics for German were computed similarly, and we allowed for inflected forms (e.g. “meiner Mutter” and “meiner Mutti” in addition to “meine Mutter” and “meine Mutti”).

A kinship term such as “mother” refers to a single individual in the family tree, but terms such as “grandmother” refer to multiple individuals in the tree. The weights in Figure 3C were computed after counts for each term were divided equally among all representatives in the tree: for example, the counts for “grandmother” were divided evenly among the two grandmothers, and the counts for “aunt” were divided equally among the four aunts.

9 Adjusting for a stable population

The analyses in the main text assume that there is a single relative for each kin type shown in Figure 2. As a result, the family trees in Figure 2 are not completely balanced: note that Alice and Bob and their parents both have four siblings each, but that Alice’s and Bob’s descendants only have a single sibling each. If the size of the population is relatively stable, then each individual should be expected to have a single sibling on average. We explored the stable-population constraint by running an analysis in which individuals were expected to have fewer

siblings, aunts, uncles, nephews, and nieces than suggested by Figure 2.

The first step is to adjust the weights in Figure 3C by assuming that the corpus statistics in Figure S6 are generated from populations where each individual tends to have a single sibling. We therefore assume that the stable-population constraint applies to the English and German populations that generated these statistics. The new set of weights includes weights for all the kin types represented in Figure 3C, but these weights now correspond to “possible weights” for relatives that may or may not exist for a given individual. We previously divided the counts in Figure S6 among the relatives represented in Figure 3C: in particular, the weights assigned to the two sisters were equal to one half of the total counts for “my sister,” and the weights assigned to the four aunts were equal to one quarter of the total counts for “my aunt.” For the stable population analysis, the possible weights assigned to each sister and each brother were twice the total counts for “my sister” and “my brother,” because each speaker has 0.5 sisters and 0.5 brothers on average, and the raw corpus statistics therefore underweight the frequency with which speakers who have sisters refer to these sisters. The possible weights assigned to each aunt, uncle, niece, and nephew were equal to the total counts for “my aunt,” “my uncle,” “my niece” and “my nephew,” because each speaker has a single aunt, uncle, niece, and nephew on average. All other possible weights were identical to the values shown in Figure 3C.

The second step is based on the assumption that the stable-population constraint applies to the cultures included in the Murdock data set. We allow for the fact that different individuals in these populations may have family trees of different sizes, and work with a probability distribution over these trees which ensures that each individual tends to have one sibling on average. For simplicity, we restricted our analysis to family trees that are subsets of Bob’s family tree in Figure 2. Each aunt, uncle, and sibling in the tree was assigned an “existence probability” of 0.25, and we enumerated the 4096 distinct trees that can be generated by tossing coins with a bias of 0.25 to determine whether each aunt, uncle or sibling exists. If a given sibling exists,

that sibling is assumed to have two children, but children of non-existent siblings are assumed not to exist. Each tree leads to a different set of need probabilities, and the need probabilities for a given tree are computed by setting the possible weight for each nonexistent relative in Figure 3C to zero and renormalizing so that the weights for all remaining relatives sum to 1. We compute the communicative cost for a given classification system by integrating over all possible sets of need probabilities.

After making the adjustments just described for a stable population, we recomputed the communicative costs for the 410 distinct kin classification systems in the data set. The correlation between the original and the revised costs was high ($r = 0.94$), which suggests that adjusting for a stable population has little impact on our results.

10 Analyses from main text

Now that we have introduced our general approach to generating and scoring possible kin classification systems, we provide more details about the specific analyses in Figures 4 and 5.

10.1 Figures 4A and 4B: partitions of the family tree

The grey points in Figure 4B represent all 831,732,650 possible kin classification systems that can be constructed using the 121 attested kinship categories that appear more than twice in our data set. As already described, we enumerated these systems using the Dancing Links algorithm.

The grey points in Figure 4A include 4×10^8 partitions sampled from the full space of partitions already described. We found, however, that most of these partitions scored substantially worse than the attested partitions, and supplemented the set of samples with two additional groups of partitions to ensure that we had included the best-scoring candidates. The first group

includes 4×10^8 partitions sampled from all conjunctive partitions, or partitions that can be represented without using a disjunction. The second group includes all of the partitions plotted in Figure 4B.

Figure 4 suggests that attested systems tend to achieve a near-optimal tradeoff between complexity and communicative cost. By comparing attested systems with randomly-chosen alternatives from Figure 4B that are matched for complexity or cost, we can explore whether attested systems tend to be near-optimal with respect to each dimension considered in isolation. The results in Figure S7 suggest that attested systems (black cross) tend to be lower in complexity than alternatives that are matched in cost, and lower in cost than systems that are matched in complexity. The strongest test along these lines is the shuffling test described earlier: we can compare the Murdock data with randomly shuffled data sets that include exactly the same categories with exactly the same frequencies but that organize these categories into different systems. Because the real and the shuffled data sets include exactly the same categories, shuffling the data leaves the average communicative cost unchanged, but Figure S7 shows that the average complexity is lower for the real data set. One thousand alternative data sets were sampled for each analysis in Figure S7, and the score for the real data set was superior to each of these alternatives, indicating that the result of each analysis is significant at the $p < 0.001$ level.

10.2 Figure 4C: grandparents, grandchildren, siblings, aunts, uncles and niblings

Each analysis in Figure 4C considers partitions of a subset of the 56 relatives in Figure 2. For each analysis we started with the full set of 71,000 ego-relative extensions, dropped all “external ego-relative extensions,” or extensions that included at least one relative from outside the current subset, and used the Dancing Links algorithm to enumerate all partitions that could

be constructed from the ego-relative extensions that remained. Even though we removed all external ego-relative extensions before enumerating the partitions for each analysis, note that external categories are typically required in order to define the categories in a given partition. For example, some partitions of the eight grandparents are defined using the categories `mother` and `father` even though these categories are not included in the actual partitions under consideration.

The cross-cultural data set used for Figures 4A and 4B included 478 partitions, but each analysis in Figure 4C considered more than 478 partitions. For instance, if the information available for a given language was missing a code for aunts but was otherwise complete, we dropped this language for the analyses in Figures 4A and 4B but included it when analyzing terms for grandparents, grandchildren, siblings, uncles, and niblings.

Figure 4C suggests that frequent systems (large black circles) tend to lie closer to the optimal frontier than rare systems (small black circles). Let the *dominance rank* of a system be the number of alternatives by which it is dominated—for example, systems on the optimal frontier will have a dominance rank of zero, and in general smaller dominance ranks indicate more dominant systems. Figure S8 plots the dominance rank of each system in Figure 4C against its frequency and confirms that the more frequent systems tend to have relatively small dominance ranks.

Tables S1, S2, S3, S4, S5 and S6 show the attested systems for each subtree (i.e. the systems that correspond to the black circles in Figure 4C). The tables include all attested systems that appear twice or more and all possible systems with a dominance rank of zero. Tables S1 through S6 demonstrate some of the strengths and limitations of our theory. Each table is sorted by frequency, and the high-frequency systems near the top of each table tend to have small dominance ranks. In all cases except Table S4, however, there are systems which have dominance ranks of zero but are never found in practice. Inspecting these “missing systems” reveals some aspects

of kin classification that our approach does not capture. For example, six of the seven missing grandparent and grandchildren systems rely on the primitives $\text{SAMESEX}(\cdot, \cdot)$ and $\text{DIFFSEX}(\cdot, \cdot)$. Social factors external to our account could explain why these primitives play a role in sibling classification but are almost never used when classifying grandparents and grandchildren—for example, the sex of an individual from the same generation might carry more social relevance than the sex of an individual two generations removed, because marriage partners are more likely to be drawn from near rather than distant generations.

10.3 Figures 5B and 5C: permutation analysis

The transformations we consider are defined in terms of the five chunks shown in Figure 5A. Note that each chunk includes two males and two females. We also consider a “flipped” version of each chunk where the two males are exchanged and the two females are exchanged. For example, the original version of the grandparents chunk is [MM, MF, FM, FF], and the flipped version is [FM, FF, MM, MF]. The five chunks can be permuted in 120 ways, and allowing each chunk to be flipped produces a total of $120 \times 2^5 = 3840$ permutations.

For any given system, some permutations may seem more natural than others. English, for example, includes two categories—aunt and uncle—that include individuals from more than one chunk. As a result, permutations of English that move the maternal siblings or the paternal siblings will violate category boundaries and are therefore likely to produce unnatural results. For instance, permuting the English system in Figure 2A by exchanging maternal siblings and grandparents produces a transformed system where Alice uses the same term for granddaughters and paternal aunts. Because of this unnaturalness, we considered only permutations that do not violate category boundaries.

In principle, each of the 3840 permutations could be applied to each of the 478 systems in the Murdock data. 35% of these possibilities remain after we remove all permutations that violate

category boundaries, as in the English aunt/uncle example above, and the results in Figure 5 are computed over this remaining set. Our analysis is therefore based on the permutations that seem most natural— because we predict that transformations tend to score poorly in general, a conservative test of our prediction should focus on the permutations with the best chance of scoring well.

Figure 5C shows results for three specific permutations. In each case, the permutation is only applied to systems where it does not violate category boundaries. As a result, the total number of systems analyzed is different across the three panels.

10.4 Figure 5D: conjunctive vs disjunctive categories

Consider two sets of rules: a disjunctive set, which includes all ten rules in Figure 3B, and a conjunctive set, which includes all rules except the three disjunctions. A category will be called *conjunctive* if it can be defined using the conjunctive set, and *disjunctive* if it can be defined using the disjunctive but not the conjunctive set. We define the dominance rank of a category as the smallest (i.e. best) dominance rank for any partition that includes the category. For each of the five subtrees in Figure 4C, Figure 5D shows that conjunctive categories tend to have smaller dominance ranks than disjunctive categories. We can therefore conclude that conjunctive categories are more likely than disjunctive categories to contribute to near-optimal kin classification systems.

Because conjunctive categories are preferred by our theory, it is natural to ask whether the near-optimal distributions in Figure 4 emerge as a result of comparing attested systems to a large collection of low-scoring disjunctive systems. Figure S9 is a counterpart of Figure 4 that is constructed in exactly the same way, except that we use the conjunctive set of rules rather than the disjunctive set. Figure S9 shows that attested systems tend to occupy the optimal region of the space of possible systems even if this space is based on conjunctive categories only.

References

60. G. P. Murdock, Kin term patterns and their distribution, *World Cultures* **11**, 102 (2000).
61. C. Goddard, Lexico-semantic universals: A critical overview, *Linguistic Typology* **5**, 1 (2001).
62. D. L. Everett, Cultural constraints on grammar and cognition in Pirahã: Another look at the design features of human language, *Curr. Anthropol.* **46**, 621 (2005).
63. A. Macfarlane, Analysis of relationships of consanguinity and affinity, *J. Anthropol. Inst. Great Britain Ireland* **12**, 46 (1883).
64. D. Knuth, Dancing links, in *Millennial Perspectives in Computer Science*, J. Davies, B. Roscoe, J. Woodcock, Eds. (Palgrave, Basingstoke, 2000), pp. 187–214.




Figure S1: Family trees for the cousins analysis.

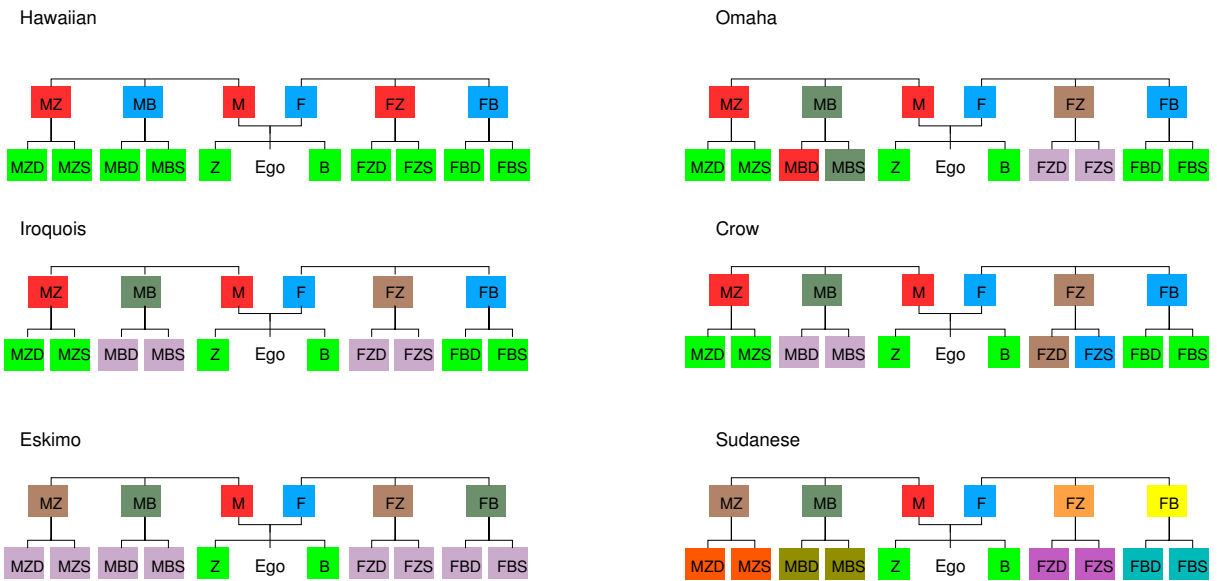


Figure S2: Six families of cousin systems.




Figure S3: Communicative cost versus complexity for a large space of possible cousin systems. Attested cousin systems are shown as black circles.

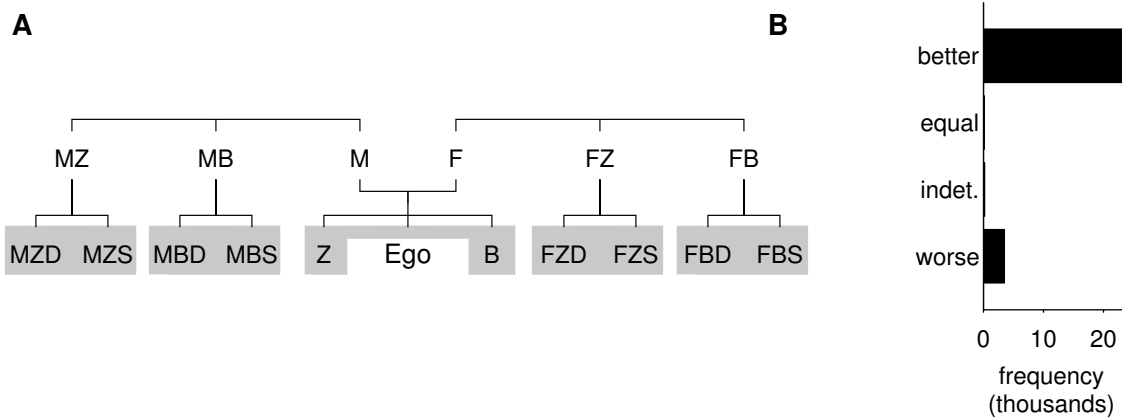


Figure S4: Fine-grained optimality analysis of attested cousin systems. (A) The gray bars indicate the five chunks used for the permutation analysis. (B) Results of the permutation analysis. Attested systems typically dominate permuted versions of these systems.

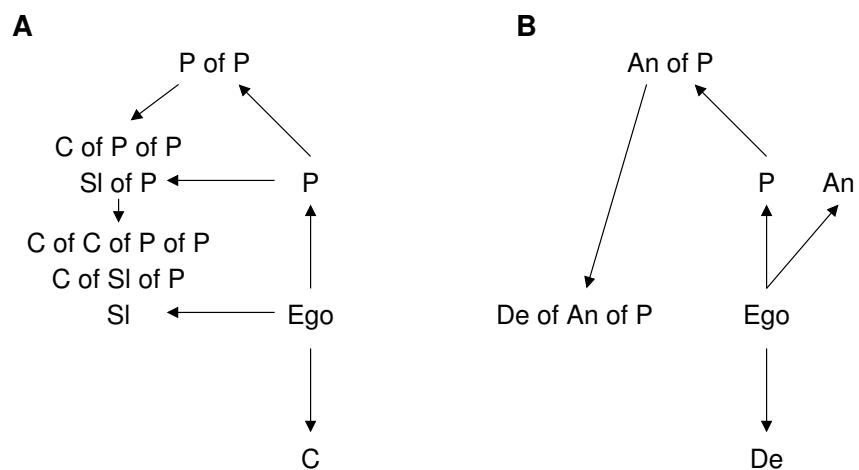


Figure S5: Procedure for generating the set of categories \mathcal{C} used for the cousins analysis. Groups P, C, Sl, An and De include variants of “parent,” “child,” “sibling,” “ancestor” and “descendant” respectively. **(A)** Combinations involving groups P, C and Sl. **(B)** Combinations involving groups An and De.




Figure S6: Corpus statistics for English and German.




Figure S7: Average cost and complexity for the Murdock data (black cross) compared to shuffled versions of the data and alternative data sets that are matched for cost or complexity. All error bars show 99% prediction intervals.

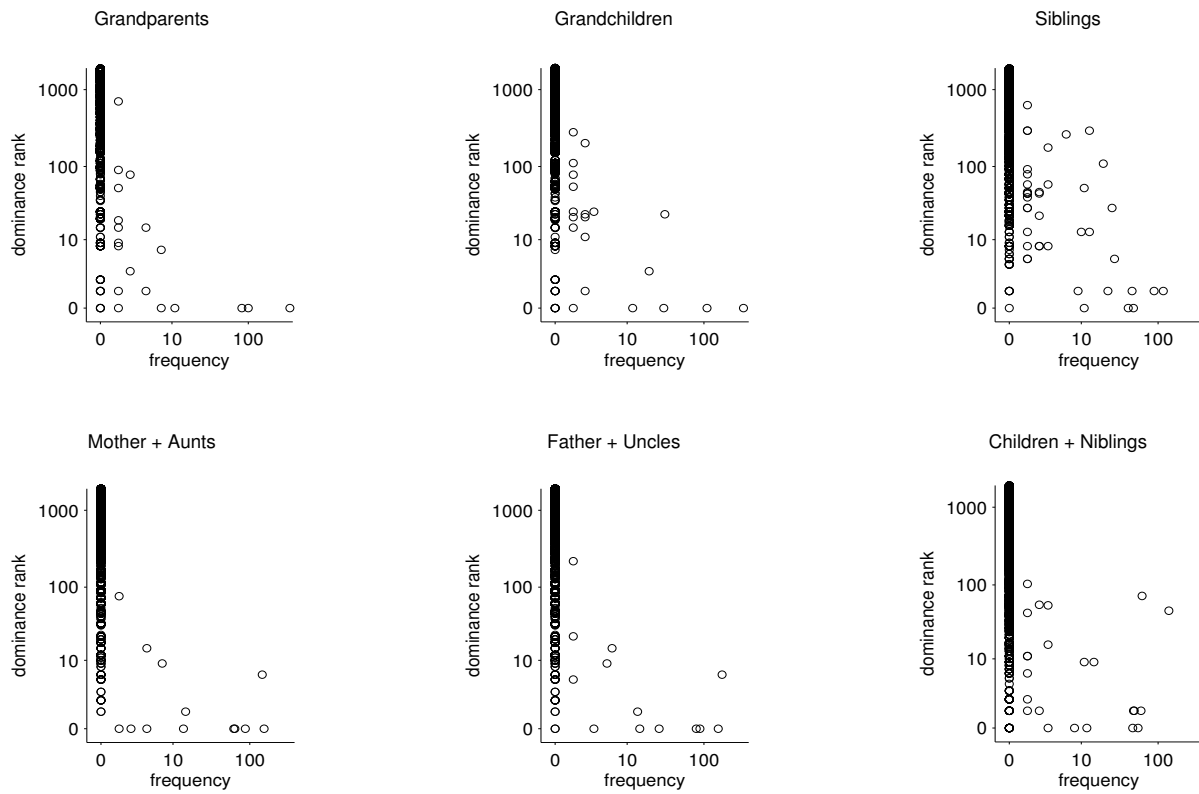


Figure S8: Dominance ranks versus frequency for six subsets of the full family tree in Figure 2. The systems that appear most frequently tend to be dominated by relatively few alternatives.




Figure S9: A version of Figure 4 based on kin categories with conjunctive definitions.

Alice (♀)				Bob (♂)				complexity	cost	freq	dom rank
MM	MF	FM	FF	MM	MF	FM	FF				
1	2	1	2	1	2	1	2	3	1	331	0
1	1	1	1	1	1	1	1	1	1.99	100	0
1	2	3	4	1	2	3	4	6	0	83	0
1	2	3	2	1	2	3	2	5	0.43	11	0
1	2	1	3	1	2	1	3	5	0.57	7	7
1	1	2	2	1	1	2	2	4	0.99	7	0
1	2	3	3	1	2	3	3	5	0.49	4	1
1	2	1	2	1	3	1	3	5	1	4	15
1	2	1	2	3	4	3	4	7	1	2	78
1	1	1	1	2	2	2	2	3	1.99	2	3
1	2	1	2	2	1	2	1	3	1	1	0
1	2	3	4	4	3	2	1	6	0	0	0
1	2	2	1	1	2	2	1	4	0.99	0	0
1	1	2	2	2	2	1	1	4	0.99	0	0

Table S1: Kin terms for grandparents. Each row represents a kin classification system—for example, the first row represents a system where Alice and Bob use the same two terms (one for “grandmother” and the other for “grandfather”). The final two columns show the empirical frequencies and dominance ranks for each system. Systems below the line have a dominance rank of zero but do not appear in the Murdock data.

	Alice (♀)				Bob (♂)							
DD	DS	SD	SS	DD	DS	SD	SS	complexity	cost	freq	dom	rank
1	1	1	1	1	1	1	1	1	1.96	310	0	
1	2	1	2	1	2	1	2	3	1	109	0	
1	1	2	2	3	3	4	4	6	0.96	32	23	
1	2	3	4	1	2	3	4	6	0	31	0	
1	1	1	1	2	2	2	2	3	1.96	20	3	
1	1	2	2	1	1	2	2	4	0.96	12	0	
1	1	1	1	2	3	2	3	5	1.48	3	25	
1	2	3	3	1	2	3	3	5	0.48	2	1	
1	1	2	2	3	3	2	2	5	0.96	2	11	
1	1	2	2	3	3	3	3	5	1.46	2	21	
1	1	2	2	2	2	3	3	6	0.96	2	23	
1	2	3	4	5	6	7	8	14	0	2	202	
1	2	1	2	2	1	2	1	3	1	1	0	
1	2	3	4	4	3	2	1	6	0	0	0	
1	2	2	1	1	2	2	1	4	0.96	0	0	
1	1	2	2	2	2	1	1	4	0.96	0	0	
1	2	1	3	1	2	1	3	5	0.39	0	0	

Table S2: Kin terms for grandchildren.

Alice (φ)				Bob (σ)				complexity	cost	freq	dom rank
Zy	Ze	By	Be	Zy	Ze	By	Be				
1	2	3	4	1	2	3	4	7	0	116	1
1	1	2	2	1	1	2	2	3	1	89	1
1	2	1	2	1	2	1	2	3	0.99	49	0
1	2	1	3	1	2	1	3	5	0.49	47	1
1	1	1	1	1	1	1	1	1	1.99	42	0
1	2	3	3	3	3	1	2	5	0.5	28	5
1	2	3	3	4	4	1	2	7	0.5	26	28
1	1	2	2	2	2	1	1	3	1	23	1
1	1	2	2	3	3	4	4	7	1	20	109
1	1	2	2	3	3	1	1	5	1	13	13
1	2	3	3	4	4	5	6	10	0.5	13	293
1	1	2	3	1	1	2	3	5	0.44	11	0
1	2	3	4	5	6	7	8	12	0	11	52
1	1	2	2	2	2	3	3	5	1	10	13
1	2	3	2	1	2	3	2	5	0.49	9	1
1	1	2	2	3	3	4	5	9	0.72	6	263
1	2	3	4	1	2	3	5	9	0	3	8
1	2	3	3	3	3	4	5	8	0.5	3	58
1	2	3	3	4	4	1	5	9	0.5	3	177
1	2	3	4	5	6	1	2	9	0	2	8
1	2	3	4	1	2	5	4	9	0	2	8
1	2	3	4	5	5	3	4	9	0.22	2	22
1	2	3	4	1	2	1	4	9	0.25	2	44
1	1	2	2	2	2	3	4	7	0.72	2	46
1	2	3	4	3	4	1	2	6	0	0	0

Table S3: Kin terms for siblings.

Alice (φ)					Bob (σ)					complexity	cost	freq	dom rank
MZy	MZe	M	FZy	Fze	MZy	MZe	M	FZy	FZe				
1	1	2	3	3	1	1	2	3	3	6	0.04	151	0
1	1	1	2	2	1	1	1	2	2	7	0.19	143	6
1	1	2	1	1	1	1	2	1	1	4	0.08	88	0
1	1	1	1	1	1	1	1	1	1	3	0.33	65	0
1	2	3	4	4	1	2	3	4	4	9	0.02	63	0
1	2	3	1	2	1	2	3	1	2	7	0.04	15	1
1	2	3	4	5	1	2	3	4	5	10	0	14	0
1	1	2	3	3	1	1	2	4	4	9	0.04	7	9
1	2	3	4	2	1	2	3	4	2	9	0.02	4	0
1	1	1	2	2	1	1	1	3	3	8	0.19	4	15
1	2	3	1	4	1	2	3	1	4	9	0.02	2	0
1	1	2	3	4	1	1	2	3	4	9	0.02	1	0

Table S4: Kin terms for aunts

Alice (φ)					Bob (σ)					complexity	cost	freq	dom rank
MBy	MBe	F	FBy	FBe	MBy	MBe	F	FBy	FBe				
1	1	2	2	2	1	1	2	2	2	7	0.26	168	6
1	1	2	3	3	1	1	2	3	3	6	0.06	151	0
1	1	2	3	4	1	1	2	3	4	9	0.03	89	0
1	1	2	1	1	1	1	2	1	1	4	0.12	80	0
1	1	1	1	1	1	1	1	1	1	3	0.45	27	0
1	2	3	4	5	1	2	3	4	5	10	0	15	0
1	2	3	1	2	1	2	3	1	2	7	0.06	14	1
1	1	2	2	2	3	3	2	2	2	8	0.26	6	15
1	1	2	3	3	4	4	2	3	3	9	0.06	5	9
1	2	3	4	2	1	2	3	4	2	9	0.03	3	0
1	2	3	4	4	1	2	3	4	4	9	0.03	0	0
1	2	3	1	4	1	2	3	1	4	9	0.03	0	0

Table S5: Kin terms for uncles.

				Bob (♂)						complexity	cost	freq	dom rank
ZyD	ZyS	ZeD	ZeS	D	S	ByD	ByS	BeD	BeS				
1	1	1	1	2	3	2	3	2	3	8	0.16	136	46
1	2	1	2	3	4	3	4	3	4	9	0.15	63	72
1	1	1	1	2	3	4	4	4	4	6	0.06	61	1
1	1	1	1	2	3	1	1	1	1	4	0.09	56	0
1	2	1	2	3	4	1	2	1	2	5	0.06	50	1
1	2	1	2	1	2	1	2	1	2	4	0.26	49	1
1	2	1	2	3	4	5	6	5	6	8	0.03	48	0
1	1	1	1	2	3	4	4	5	5	9	0.05	15	9
1	2	1	2	3	4	5	5	5	5	7	0.05	12	0
1	1	2	2	3	4	5	5	6	6	11	0.03	11	9
1	1	1	1	2	3	4	5	4	5	7	0.05	8	0
1	2	1	2	3	4	5	6	7	8	12	0.02	3	0
1	1	1	1	2	3	4	5	6	7	11	0.03	3	16
1	2	1	2	1	3	1	3	1	3	8	0.20	3	54
1	2	1	2	3	4	1	5	1	5	8	0.04	2	1
1	2	1	2	3	2	3	2	3	2	8	0.21	2	55
1	2	3	4	5	6	7	8	9	10	16	0	0	0
1	2	3	4	5	6	7	8	7	9	15	0.007	0	0
1	2	3	4	5	6	7	8	7	8	12	0.02	0	0
1	2	3	4	5	6	2	1	4	3	9	0.03	0	0
1	2	3	4	5	6	1	7	3	8	13	0.01	0	0
1	2	1	3	4	5	6	7	8	9	15	0.007	0	0
1	2	1	3	4	5	6	7	6	7	11	0.02	0	0
1	2	1	2	3	4	5	6	5	7	11	0.02	0	0
1	2	1	2	3	4	2	5	2	5	7	0.05	0	0
1	2	1	2	3	4	5	1	5	1	7	0.05	0	0
1	2	1	2	3	4	2	1	2	1	5	0.06	0	0
1	1	1	1	2	2	1	1	1	1	2	1.06	0	0

Table S6: Kin terms for nieces, nephews and children (male speaker).