A general principle governing neuronal evolution reveals a human-accelerated neuron type potentially underlying the high prevalence of autism in humans

Alexander L. Starr*, Hunter B. Fraser*

Department of Biology, Stanford University, Stanford, CA 94305, USA

*Corresponding authors: astarr97@stanford.edu, hbfraser@stanford.edu

Abstract

The remarkable ability of a single genome sequence to encode a diverse collection of distinct cell types, including the thousands of cell types found in the mammalian brain, is a key characteristic of multicellular life. While it has been observed that some cell types are far more evolutionarily conserved than others, the factors driving these differences in evolutionary rate remain unknown. Here, we hypothesized that highly abundant neuronal cell types may be under greater selective constraint than rarer neuronal types, leading to variation in their rates of evolution. To test this, we leveraged recently published cross-species single-nucleus RNA-sequencing datasets from three distinct regions of the mammalian neocortex. We found a strikingly consistent relationship where more abundant neuronal subtypes show greater gene expression conservation between species, which replicated across three independent datasets covering $>10^6$ neurons from six species. Based on this principle, we discovered that the most abundant type of neocortical neurons—layer 2/3 intratelencephalic excitatory neurons—has evolved exceptionally quickly in the human lineage compared to other apes. Surprisingly, this accelerated evolution was accompanied by the dramatic down-regulation of autism-associated genes, which was likely driven by polygenic positive selection specific to the human lineage. In sum, we introduce a general principle governing neuronal evolution and suggest that the

exceptionally high prevalence of autism in humans may be a direct result of natural selection for lower expression of a suite of genes that conferred a fitness benefit to our ancestors while also rendering an abundant class of neurons more sensitive to perturbation.

Introduction

With the advent of single cell RNA-sequencing (scRNA-seq), it became possible to systematically delineate molecularly defined cell types across the brain[1,2]. As more large-scale datasets were published, it quickly became clear that the mammalian brain contains a staggering array of neuronal cell types, with recent whole-brain studies identifying nearly as many neuronal types as there are protein-coding genes in the genome[1–3]. In addition, cross-species atlases in the neocortex revealed that most cortical neuronal types are highly conserved in primates and rodents, with very few neocortical neuronal types being specific to primates and none being entirely specific to humans[4–8]. This suggests that divergence involving homologous cell types—such as their patterns of gene expression, relative proportions, and connectivity—may play a central role in establishing uniquely human cognition.

Two decades before the generation of these cross-species cell type atlases, the first whole-genome sequences of eukaryotes were published, enabling genome-wide studies of evolution for the first time[9]. One of the first questions to be addressed in the nascent field of evolutionary genomics was why some proteins are highly conserved throughout the tree of life, whereas others evolve so quickly as to be almost unrecognizable as orthologs even over relatively short divergence times[10–13]. A protein's expression level emerged as the strongest and most universal predictor of its evolutionary rate, with highly expressed proteins accumulating fewer protein-coding changes due to greater constraint[10,14–16].

In contrast to tens of thousands of publications about the evolutionary rates of proteins[17], the evolutionary rates of cell types, another key building block of multicellular life, have received relatively little attention[18]. Just as different proteins make up every cell, different cell types make up every multicellular organism. Furthermore, just as protein evolutionary rates are measured by the total rate of change of their amino acids, the evolutionary rates of cell types—which are typically defined by their patterns of gene expression—can be measured by divergence in genome-wide gene expression[4–8]. For example, it is well-established that gene expression in neurons is more conserved between humans and mice than gene expression in glial cell types such as astrocytes, oligodendrocytes, and microglia[19]. Previous analogies between genes and neural cell types have been fruitful for understanding the evolution of novel cell types[6,20–23], providing an encouraging precedent for our analogy.

One area that has been explored more thoroughly is the association of specific cell types with human diseases and disorders[24]. For example, integration of gene-trait associations with cell type-specific expression profiles has revealed that microglia likely play a central role in Alzheimer's disease[25,26]. Similar analyses have also revealed that layer 2/3 intratelencephalic excitatory (L2/3 IT neurons)—which enable communication between neocortical areas[27] and are thought to be important for uniquely human cognitive abilities[27,28]—likely play a particularly important role in autism spectrum disorder (ASD) and schizophrenia (SCZ)[29–36], together with deep layer IT neurons[36–38]. ASD and SCZ are neurodevelopmental disorders with different but overlapping characteristics, including major effects on social behavior[39–41]. Interestingly, individuals with ASD are more likely to be diagnosed with SCZ than individuals without an ASD diagnosis[39,42–44]. Furthermore, there is a strong overlap in the genes that have been implicated in both disorders[36,39].

From an evolutionary perspective, it has been proposed that ASD and SCZ may be unique to humans[45–47]. This is primarily based on two main lines of reasoning. First, ASD- and SCZ-

associated behaviors that could reasonably be observed in non-human primates (e.g. SCZ-associated psychosis) have been observed either infrequently or not at all in non-human primates[46]. However, ASD-like behavior has been observed in non-human primates[48] and the difficulties inherent to cross-species behavioral comparisons combined with relatively low sample sizes make it difficult to compare the prevalence of these behaviors in human and non-human primate populations. Second, core ASD- and SCZ-associated behavioral differences involve cognitive traits that are either unique to or greatly expanded in humans (e.g. speech production and comprehension or theory of mind)[49–53]. As a result, certain aspects of ASD and SCZ are inherently unique to humans.

While comparing interindividual behavioral differences across species remains challenging, recent molecular and connectomic evidence lend credence to the idea that the incidence of ASD and SCZ increased during human evolution. For example, large-scale sequencing studies in both ASD and SCZ cohorts have identified an excess of genetic variants in human accelerated regions (HARs)—genomic elements that were largely conserved throughout mammalian evolution but evolved rapidly in the human lineage[54–56]. Furthermore, transcriptomic studies have identified a human-specific shift in the expression of some synaptic genes during development that is disrupted in ASD[57]. In addition, connectomic studies have shown that human-chimpanzee divergence in brain connectivity overlaps strongly with differences between humans with and without SCZ[58]. Overall, evidence suggests that ASD and SCZ may be particularly prevalent in humans, but the factors underlying this increased prevalence remain unknown. Positive selection—also known as adaptive evolution—of brain-related traits in the human lineage has been proposed to underlie this increase[45–47,59,60]. Although this idea is supported by the links between HARs (many of which are thought to have been positively selected[56]) and ASD and SCZ, there is no direct evidence for positive selection on the expression of genes linked to ASD and SCZ.

Here, we set out to test whether the inverse relationship between abundance and evolutionary rates—which has been well-established for proteins[10,14–16]—might also hold for cell types. We found a robust negative correlation between cell type proportion and evolutionary divergence in the neocortex, suggesting that this relationship holds at multiple levels of biological organization. Based on this, we identify unexpectedly rapid evolution of L2/3 IT neurons and strong evidence for polygenic positive selection for reduced expression of ASD-linked genes in the human lineage, suggesting that positive selection may have increased the prevalence of ASD in modern humans.

Results

Cell type proportion as a general factor governing the rate of neuronal evolution

Based on the gene-cell type analogy outlined above, we hypothesized that a change in gene expression in a more abundant cell type may tend to have more negative fitness effects than the same change in a less abundant cell type (Figure 1A). If this were the case, this would lead to greater selective constraint, and thus slower divergence, of global gene expression in more abundant cell types.

Testing this hypothesis requires comparing two quantities: cell type proportions and the evolutionary divergence in genome-wide gene expression levels between orthologous cell types across species. Importantly, both quantities can be estimated from the same single-nucleus RNA-seq (snRNA-seq) data, facilitating comparison between them. To ensure sufficient statistical power, we searched the literature for published snRNA-seq data sets that fulfilled a stringent pair of criteria. First, they must have multiple species profiled in the same study using

the same snRNA-seq protocols for each species within a study. Second, they must contain at least 10 orthologous cell types having 250 or more cells per species (not including immune cells, as these do not have stable cell type proportions). We identified three studies fulfilling these criteria, focused on three distinct regions of the mammalian neocortex: medial temporal gyrus (MTG), dorsolateral prefrontal cortex (DLPFC), and primary motor cortex (M1)[5,7,8]. All three studies included samples from 3-5 species, including human and marmoset, with 300,000 – 500,000 neuronal nuclei profiled per study[5,7,8]. These nuclei were clustered into between 12 – 17 neuronal subclasses (with at least 250 cells per species) in each study, which we then used for our analyses[5,7,8]. Throughout, we use the term cell type for the general concept of different types of cells and as an umbrella term for both subclasses and subtypes, use the term subclass for the traditional classification of neuronal types found in the neocortex, and reserve the term subtype for more fine-grained clustering of cells.

To test our hypothesis, we began by comparing human and marmoset (the only pair of species present in all three datasets) in the MTG, which had the greatest sequencing depth. We first estimated gene expression divergence for each of 14 subclasses using the Spearman correlation distance (1 – Spearman's rho) between the pseudobulked expression of each species for each neuron subclass, restricting to one-to-one orthologous genes (see Methods). We observed a surprisingly strong negative correlation between subclass proportion and gene expression divergence (Spearman's rho = -0.84, $p = 8.0 \times 10^{-5}$, Figure 1B), indicating that more abundant neuronal subclasses showed greater conservation of genome-wide gene expression. To ensure that estimates of cell type-specific expression divergence were not biased by cell type proportion itself, we analyzed the same number of cells and total reads for each cell type in each species. Specifically, for all analyses we report the median rho and p-values from 100 independent down-samplings of cells and pseudobulked counts without replacement (see Methods).

We next asked whether the same pattern was present in the other cortical regions. We observed a similar strong negative correlation in the two other independently generated datasets (Spearman's rho = -0.76, p = 0.00041 in the DLPFC, Figure 1C; Spearman's rho = -0.73, p = 0.0065 in the M1, Figure 1D). This replication suggests that the relationship we observed holds true across the primate neocortex. In addition, the fact that methodological details and biological samples differ across these studies lends additional robustness to any patterns shared by all three.

To explore the generality of this result in additional species, we repeated this analysis between every pair of species in each dataset. We observed similarly strong negative correlations across all pairwise comparisons (Supplemental Figures 1-3), with the interesting exception of comparisons between humans and non-human great apes, where a weaker negative correlation was observed (discussed below). Furthermore, we observed strong negative correlations within excitatory or inhibitory subclasses in all three brain regions (Figure 2 and Supplemental Figures 4-9, although this correlation does not reach statistical significance for inhibitory neurons in M1, potentially due to having only five subclasses in that dataset). In addition, we tested all possible combinations of a wide variety of filtering parameters, analysis decisions, and distance metrics, finding that this negative correlation was generally robust to any reasonable choice of parameters we made (Supplemental Table 1).

Next, we investigated this relationship at the level of neuronal subtypes, a finer-grained clustering with ~4-fold more cell subtypes than subclasses. We found strong negative correlations between subtype proportion and expression divergence when using all neurons (Figure 3A-C, Supplemental Figures 10-12) or only excitatory neurons (Figure 3D-F, Supplemental Figures 13-15). When restricting our analysis to inhibitory neurons, this

correlation was statistically significant in the MTG and in two of three comparisons (mouse-marmoset and human-mouse) in the M1, but not in DLPFC (Figure 3G-I, Supplemental Figures 16-18). This may reflect the lower read depth (average of 180,054 counts used for DLPFC, compared to 254,703 for M1 and 325,422 for MTG) or lower numbers of cells per subtype in the DLPFC data compared to the other datasets, as we observed a much stronger negative correlation (Spearman's rho = -0.50, p = 0.057) when restricting to subtypes with at least 500 cells in the DLPFC data (Supplemental Figure 19). Overall, our results suggest that there is a strong, robust negative correlation between expression divergence and cell type proportion for neocortical neurons.

Finally, we investigated the properties of the genes driving the negative correlation we observed. First, we stratified genes into three equally sized bins by their expression level and recomputed correlations in each bin. Interestingly, while we observed strong correlations for highly and moderately expressed genes, there was no significant correlation when restricting to lowly expressed genes (Figure 4A, Supplemental Figures 20-22, Supplemental Table 2). Next, we stratified genes based on evolutionary constraint on expression level or cell type-specificity of expression (using s_{het}[61] and the Tau metric[62] respectively, Supplemental Tables 3 and 4). While there was no difference in correlation when stratifying by constraint on expression (Supplemental Figures 23-25, Supplemental Table 3), we observed a much stronger negative correlation between cell type proportion and expression divergence for more cell type-specifically expressed genes (Figure 4B, Supplemental Figures 26-28, Supplemental Table 4). Since expression level is also associated with cell-type specificity, we tested whether these two properties were contributing independently to the negative correlations by stratifying genes by one of them while simultaneously controlling for the other. We found that both properties retained their predictive power even when controlling for the other (Figure 4C-D, Supplemental Figures 29-34, Supplemental Tables 2 and 4), suggesting independent contributions. We note

that whether the weaker correlations we observed for lowly expressed genes were due to a true lack of association or simply less accurate expression level measurements remains an open question that will require larger datasets to explore. Overall, our results suggest that more highly expressed, cell type-specific genes are primarily driving the negative correlation between cell type proportion and gene expression divergence.

Rapid evolution of layer 2/3 intratelencephalic neurons in the human lineage

Having identified this strong relationship between cell type proportion and evolutionary divergence, we reasoned that cell types with much faster divergence in the human lineage than expected based on their abundance may have been subject to atypical selective forces.

To identify subclasses showing the most dramatic lineage-specific shifts in selection, we decomposed human-chimpanzee MTG expression divergence into its two components, divergence on the human branch and divergence on the chimpanzee branch. Applying the concept of parsimony—explaining the data with as few evolutionary transitions as possible—allows an outgroup species such as gorilla to polarize changes and assign them to either the human or chimpanzee branch (see Methods). In the chimpanzee lineage, there was a strong negative correlation between divergence and subclass proportion (Figure 5A, Spearman's rho = -0.77, p = 0.00076), similar to the correlations between other primate species (Figure 1A, Supplemental Figure 1). However, we observed a much weaker negative correlation in the human lineage (Figure 5B, Spearman's rho = -0.19, p = 0.49). The clearest outlier weakening the correlation was L2/3 IT neurons, the most abundant neuronal subclass, which diverged much faster than expected based on its proportion. This was also true to a lesser extent for the next two most abundant subclasses, L4 IT and L5 IT neurons. Indeed, removing these three subclasses substantially strengthened the negative correlation between subclass proportion and human-specific divergence (Figure 5B; Spearman's rho = -0.59, p = 0.041), making it

indistinguishable from the corresponding chimpanzee-specific correlation (Figure 5A, blue points; Spearman's rho = -0.58, p = 0.048). Quantifying the magnitude of human acceleration for every subclass confirmed that L2/3 IT neurons underwent the greatest acceleration, followed by L4 and L5 IT neurons (Figure 5C).

Accelerated evolution can involve either positive selection favoring gene expression changes that increased fitness, or relaxed selective constraint in which random mutations are allowed to accumulate over time because they have little or no effect on fitness[56]. Although both positive selection and relaxed constraint can lead to similar patterns of lineage-specific acceleration, they imply very different underlying factors: positive selection is the force underlying nearly all evolutionary adaptation, while relaxed constraint is simply the weakening or absence of natural selection which can lead to the passive deterioration of genes and their regulatory elements via mutation accumulation.

To distinguish whether positive selection or relaxed constraint was more likely to underlie the human-specific acceleration of IT neurons, we investigated the interindividual variability in expression of each neuronal subclass in the human population[63]. If IT neurons evolved under reduced constraint in the human lineage then we would expect them to have more variable expression among humans, leading to a weaker negative correlation between subclass proportion and interindividual variability. Instead, we observed a strong negative correlation between subclass proportion and interindividual variability in gene expression, with L2/3 IT neurons having the lowest variability of any subclass among humans (Figure 5D, Spearman's rho = -0.55, p = 0.049). Consistent with this, L2/3 IT neurons had the largest human branch divergence relative to their expression variability in modern humans (Figure 5E). Overall, these results suggest that the rapid gene expression evolution of L2/3 IT neurons in the human lineage was unlikely to be due to relaxed constraint, and instead more likely the result of positive selection (Figure 5F), though we cannot formally rule out other possible scenarios (see

Discussion). In addition, it suggests that the relationship between cell type proportion and expression divergence holds within species as well as between species.

Lower expression of ASD-linked genes in humans compared to chimpanzees

As discussed above, L2/3 IT neurons are thought to play a particularly important role in ASD. To investigate a potential connection between accelerated evolution of L2/3 IT neurons and the prevalence of ASD in humans, we asked whether genes previously implicated in ASD showed human-specific gene expression patterns. To begin, we asked whether differentially expressed ASD-linked genes tended to be more highly expressed in humans or in chimpanzees, testing each neuron subtype in the DLPFC and MTG datasets. Although in some types of neurons, such as L6 CT neurons, there was no significant directionality bias (Figure 6A), many subclasses showed a bias towards lower expression of ASD-linked genes in humans compared to chimpanzees (Fig 6B). Strikingly, in both datasets we observed the most significant trend towards lower expression of these genes in human L2/3 IT neurons (60 genes higher in human vs. 12 genes lower in human in DLPFC, Figure 6C, Supplemental Figure 35).

This excess of ASD-linked genes with lower expression in humans is consistent with either down-regulation in the human lineage, up-regulation in the chimpanzee lineage, or a combination of both. To distinguish between these possibilities, we again used gorilla as an outgroup to assign each gene's expression divergence in the MTG to either the human or chimpanzee lineage.

Comparing the expression of ASD-linked genes in all three species revealed that gorilla gene expression is generally intermediate between human and chimpanzee, but closer to chimpanzee. Specifically, the distribution of ASD-linked gene expression log-ratios of [human/gorilla] is generally negative (lower expression in humans), whereas [chimpanzee/gorilla] is generally positive (Figure 6D). Interestingly, the magnitude of the

277 [human/gorilla] divergence was generally greater than the magnitude of the [chimp/gorilla]
278 divergence, suggesting that there has been greater divergence in the human lineage
279 ([human/gorilla] median absolute $\log_2$ fold-change = 0.45, [chimp/gorilla] = 0.28, t-test p =
280 0.00036, Figure 6D). Consistent with this, a larger number of ASD-linked genes' expression
281 diverged on the human branch than expected by chance in L2/3 IT neurons (binomial p = 0.025,
282 1.5-fold enrichment, Supplemental Figure 36). Overall, these results suggest a strikingly
283 consistent pattern of human-specific down-regulation of ASD-associated genes in a neuronal
284 cell type with a key role in ASD.

285 ***Polygenic positive selection for down-regulation of ASD-linked genes in the human***
286 ***lineage***

287 This human-specific down-regulation of ASD-linked genes is striking and, based on the highly
288 constrained expression of these genes, likely functionally significant. However, as with the
289 accelerated evolution of L2/3 IT neurons discussed above (Figure 5), the question of whether
290 lineage-specific selection was responsible is key to understanding the factors that drove this
291 divergence in the human lineage. Other potential explanations fall into two main categories. One
292 is genetic changes that were not driven by selection, such as mutations that had little effect on
293 fitness but became established in the human lineage through genetic drift. The other is non-
294 genetic differences in the individuals sampled for these data sets; factors such as diet,
295 environmental exposures, and age can impact gene expression but cannot be controlled in any
296 comparison of tissue samples between humans and other species.

297 In order to definitively implicate lineage-specific selection, two steps are necessary. First, all
298 non-genetic causes must be ruled out. Although this is not possible with tissue samples, it can
299 be achieved *in vitro*. Human and chimpanzee induced pluripotent stem cells (iPSCs) can be
300 fused to generate hybrid tetraploid iPSCs, which can then be differentiated into relevant cell

types or organoids[64,65]. In each hybrid cell, the human and chimpanzee genomes share precisely the same intracellular and extracellular environment. As a result, any difference in the relative expression levels of the human and chimpanzee alleles for the same gene—known as allele-specific expression (ASE)—reflects *cis*-regulatory changes between the two alleles. Both environmental and experimental sources of variability (including batch effects) are perfectly controlled in the hybrid system, since all comparisons are between alleles that share an identical environment and are present in the same experimental samples[64,65].

The second step necessary to infer lineage-specific selection is to test, and reject, a statistical "null model" of neutral evolution for the genetic component of divergence[66]. The simplest and most robust pattern predicted under neutral evolution of gene expression is the expectation that in a comparison between two species, genetic variants causing expression divergence will be just as likely to lead to higher expression in one species as in the other[67]. For example, in a set of 20 functionally related genes, neutral evolution leads to a similar pattern as a series of 20 coin flips—an expectation of ~10 genes more highly expressed in one species and ~10 in the other, with deviation from this average following the binomial distribution[67]. In contrast, natural selection that favors lower expression of these genes in one lineage will lead to a pattern of biased expression, with most of the 20 genes expressed lower in that lineage[67]. This framework, which has been applied extensively to gene expression and other quantitative traits[64–66,68,69], is known as the sign test. Because the ASE of each gene in hybrid cells is generally independent of that of other genes, facilitating statistical analysis, hybrid ASE is ideally suited for detecting selection with the sign test whereas data from non-hybrids cannot be used in this manner.

To apply this test for lineage-specific selection, we focused on a previously published RNA-seq dataset from human-chimpanzee hybrid cortical organoids[64]. These organoids—which include glutamatergic and GABAergic neurons, astrocytes, and neural precursor cells—were sampled in a bulk RNA-seq time series spanning 200 days of development *in vitro*[64]. As described above, a

significant bias in the directionality of ASE for any predefined set of genes can reject the null hypothesis of neutral evolution, and instead suggests lineage-specific selection. Applying this test to known ASD-associated genes, we found a strong bias toward lower expression from the human allele in cortical organoids at two different stages of development (2.0-fold enrichment at day 100 of organoid development; binomial p = 0.003; Figure 6E). The bias toward lower expression from human alleles was even stronger when using only high-confidence ASD genes (2.5 fold-enrichment; binomial p = 0.01 at day 100; Supplemental Figure 37). This ASE bias is inconsistent with neutral evolution, and strongly implies the action of lineage-specific selection on the expression of ASD-linked genes.

To determine the lineage (human or chimpanzee) on which the ASD-linked gene expression changes occurred, for genes with matching directionality in the L2/3 IT and organoid data we once again polarized gene expression divergence in the MTG into human-derived and chimpanzee-derived categories using gorilla as an outgroup. Out of 17 chimpanzee-derived genes, there was no directionality bias in the organoid ASE data at either day 100 or day 150 (9 out of 17 with lower expression from the human allele at day 150, Figure 6F-G, Supplemental Figure 38), consistent with neutral evolution. However, out of 22 human-derived genes, 20 had lower expression from the human allele (Fisher's exact test p = 0.010 at day 150; odds ratio = 8.9; Figure 6F-G). This trend is even stronger when using a more relaxed false discovery rate (FDR) cutoff of 0.1 (25 down-regulated in human vs 2 up-regulated; Fisher's exact test p = 0.0017; odds ratio = 12.5). Overall, this strongly suggests that many ASD-linked genes were down-regulated specifically in the human lineage.

This coordinated down-regulation of 25 ASD-linked genes could conceivably be due to either positive selection or loss of constraint, as both of these types of lineage-specific selection could lead to down-regulation[67,69]. To determine if ASD-associated genes might be evolving under relaxed constraint in humans, we tested several predictions of the relaxed constraint model.

First, genes evolving under relaxed constraint might be expected to have accumulated more substitutions affecting protein sequence and/or gene expression in the human lineage. However, we found no difference in protein sequence constraint (measured by dN/dS[70]) or the number of mutations near the transcription start site (TSS) between humans and chimpanzees (after correcting for genome-wide differences between the two lineages, $p = 0.42$ for dN/dS, $p = 0.24$ for mutations near TSS, paired t-test, Supplemental Figure 39A-B). In addition, the expression of genes evolving under relaxed constraint in humans would likely be more variable across human individuals compared to chimpanzee individuals. However, we found the opposite for ASD-linked genes—slightly less variability in expression in humans ($p = 0.08$ for DLPFC, $p = 2.5 \times 10^{-5}$ for MTG, paired t-test, Supp Fig. 39C-D), suggesting that the expression of ASD-linked genes may actually be under stronger constraint in humans compared to chimpanzees. Consistent with this, the vast majority of ASD-linked genes have strongly constrained expression in humans as measured by loss-of-function intolerance (82% of ASD-linked genes have probability of loss of function intolerance[71] > 0.9 compared to 17% genome-wide; similarly, 82% of ASD-linked genes have a fitness effect of heterozygous loss of function[61] [s_{het}] > 0.1, compared to 18% genome-wide).

Although we cannot rule out any possibility of relaxed constraint at some point in the past, these results favor a model in which polygenic positive selection acted to decrease expression of ASD-linked genes in human L2/3 IT neurons (as well as in some other cell types in the neocortex). As loss of function underlies increased probability of ASD diagnosis for the vast majority of these genes[72], this suggests that down-regulation of ASD-linked gene expression may have increased ASD prevalence in the human lineage. In monogenic cases, decreased expression of ASD-linked genes in the human lineage may have led to humans being closer to a hypothetical “ASD expression threshold” below which ASD characteristics would manifest. As an example, *DLG4*, which encodes the key synaptic protein PSD-95 and for which loss of one

copy causes ASD[73], has 2.5-fold lower expression in humans compared to chimpanzees (Figure 6H). Consistent with this, it also has 2.5-fold lower protein abundance in the postsynaptic density (PSD) in humans compared to rhesus macaques, and 3.4-fold lower protein abundance in humans compared to mice[74] (human vs. rhesus t-test p = 0.0028, human vs. mouse t-test p = 0.00014, Supplemental Figure 40). While this human-specific down-regulation that led to the current human baseline expression level of *DLG4* is not sufficient to cause ASD, further down-regulation via loss of a single copy may push humans below the ASD expression threshold whereas loss of a single copy in chimpanzees would maintain expression above this threshold (Figure 6H). Although these genes are linked to ASD primarily due to their monogenic effects, the majority of ASD cases are thought to be caused by many small genetic and environmental perturbations collectively pushing individuals past some threshold[75]. We propose that the down-regulation of ASD-linked genes in humans increased the likelihood of ASD in the human lineage such that small perturbations on a developmental timescale are sufficient to cause ASD characteristics in humans but not chimpanzees (Figure 6I).

Down-regulation of schizophrenia-linked genes in humans

Having observed a consistent pattern of human-specific down-regulation for ASD-linked genes, we then tested whether genes linked to schizophrenia (SCZ)[76], another human-specific neuropsychiatric disorder, show a similar bias. We found an 8-fold enrichment for human down-regulation of SCZ-linked genes in DLPFC L2/3 IT neurons (Supplemental Figure 41A-B). Although this is even stronger than the ASD bias, it only reaches an FDR < 0.05 in three MTG subclasses, such as Lamp5 and Pax6 inhibitory neurons, due to much lower statistical power (31 SCZ-linked genes vs. 233 high-confidence ASD-linked). Consistent with the known genetic overlap between ASD and SCZ, six of the SCZ-linked genes are also implicated in ASD, making it difficult to disentangle the signal from ASD and SCZ. Furthermore, although there are very few SCZ-linked genes with significant ASE in the hybrid cortical organoid data, among all SCZ-

linked genes regardless of significance there is a clear bias toward human down-regulation (2.6 fold-enrichment, binomial test p = 0.025 at day 150, Supplemental Figure 41C). We interpret these results as preliminary evidence that SCZ-linked genes may have also been subject to selection for down-regulation in the human lineage, though further work will be required to confirm this.

Discussion

Building on an analogy between genes and cell types, we have identified a general principle underlying the rate of evolution of different neuronal types in the mammalian neocortex. We found a strong negative correlation between the abundance of each neuronal cell type and the rate at which its gene expression levels diverge across six mammalian species and three independent datasets[5,7,8]. Interestingly, this correlation remained very strong when collectively analyzing inhibitory and excitatory neurons, despite their very different developmental origins and functions[77,78].

Based on this initial discovery, we found that L2/3 IT neurons evolved unexpectedly quickly in the human lineage compared to other apes. This accelerated evolution included the disproportionate down-regulation of genes associated with autism spectrum disorder and schizophrenia, two neurological disorders closely linked to L2/3 IT neurons that are common in humans but rare in other apes. Finally, we found that this down-regulation, present both in adult neurons and in organoid models of the developing brain, was likely due to polygenic positive selection on *cis*-regulation. These results differ from, but do not contradict, previous findings that a group of synapse genes show human-specific up-regulation during early development that is disrupted in people with ASD[57]. Overall, our analysis suggests that natural selection on gene

426 expression may have increased the prevalence of ASD, and perhaps also SCZ, in humans (Fig
427 6H).
428
429 Although it has been widely hypothesized that natural selection for human-specific traits has
430 increased human disease risk[46,47,79–81], unambiguous evidence for this has been lacking. While
431 there is strong evidence linking natural selection on within-human genetic variation to disease
432 risk (e.g. sickle cell disease[82]), it has proven far more challenging to find similar examples
433 involving genetic variants shared by all humans. There are human-chimpanzee differences that
434 have been linked to interspecies differences in disease risk (e.g. human-specific
435 pseudogenization of the *CMAH* gene, which is thought to have shaped human susceptibility to
436 infectious diseases[81,83,84]), but there is no evidence for positive selection on these interspecies
437 genetic differences. In addition, while there are many examples of positive selection on human-
438 chimpanzee differences[64,65,70,85–87], these changes have no clear link to the likelihood of
439 diseases or disorders in humans. Finally, although the enrichment for ASD-linked variants within
440 HARs[54,55] is suggestive of a role for human-chimpanzee differences in HARs (many of which are
441 thought to be positively selected[56]) in increasing the likelihood of ASD in humans, a connection
442 between those human-chimpanzee differences and ASD has not been established. Overall, our
443 findings provide the strongest evidence to date supporting the long-standing hypothesis that
444 natural selection for human-specific traits has increased the likelihood of certain disorders.
445
446 Although our results strongly suggest natural selection for down-regulation of ASD-linked genes,
447 the reason why this conferred fitness benefits to our ancestors remains an open question.
448 Answering this question is difficult in part because we do not know what human-specific
449 features of cognition, brain anatomy, and neuronal wiring gave our ancestors a fitness
450 advantage, but we can speculate about two general classes of evolutionary scenarios. First,
451 down-regulation of ASD-linked genes may have led to uniquely human phenotypes. For

example, haploinsufficiency of many ASD-linked genes is associated with developmental delay[47], so their down-regulation could have contributed to the slower postnatal brain development in humans compared to chimpanzees. Alternatively, capacity for speech production and comprehension are unique to or greatly expanded in humans and often impacted in ASD and SCZ[53,88]. If down-regulation of ASD-linked genes conferred a fitness advantage by slowing postnatal brain development or increasing the capacity for language, that could result in the signal of positive selection we observed.

On the other hand, the down-regulation we observed may have been compensatory and reduced the negative effects of some other human-specific trait or traits. For example, the ratio of excitatory and inhibitory synapses on pyramidal neurons is fairly constant between humans and rodents despite massive differences in brain and neuron size[89]. In addition, excitatory-inhibitory imbalance is a leading hypothesis for the circuit basis of ASD[90]. If human brain expansion, changes in metabolism, or any other factor shifted this balance away from the fitness optimum, down-regulation of ASD-linked genes could potentially compensate. Overall, more work to understand human and non-human primate phenotypic differences and how polygenic changes in gene expression affect phenotypes is needed if we are to better understand selective forces acting on the expression of ASD-linked genes in the human lineage.

Our results come with important caveats. As with most correlations, causality is not implied. Our initial hypothesis was that cell type proportions may affect evolutionary rates via more severe fitness effects of expression changes in more abundant cell types, leading to greater evolutionary constraint than in rare cell types (Fig 1A). While this is a plausible explanation for our results, there also may be unknown correlates of cell type proportion that are causal. We leave explicit testing of this model to future work.

Along with establishing a mechanism underlying these correlations, another exciting future direction will be to explore this phenomenon in other tissues. This will become increasingly feasible as subclasses and fine-grained subtypes are annotated in large, uniformly processed cross-species studies. It will also be interesting to explore what factors are associated with the rate of cell type-specific gene expression divergence in contexts that lack stable cell type proportions (e.g. during development or in the immune system).

Considering that many ASD-linked genes are extremely sensitive to perturbations in their expression, our findings raise the important question of how significant reductions in the expression of so many dosage-sensitive genes were tolerated in the human lineage. As haploinsufficiency of many of these genes has severe fitness consequences in both humans and mice[47], it is unlikely that these changes occurred through single mutations of large effect. In addition, our analysis of allele-specific expression suggests that *cis*-regulatory changes underlie many of the gene expression changes we observe. Therefore, we favor a model in which many *cis*-acting mutations of small effect fixed over time, eventually leading to the large-scale down-regulation of ASD-linked genes in the human lineage. It will be interesting to use deep learning predictions of variant effects combined with experimental validation to identify the genetic differences underlying changes in the expression of ASD-linked genes in the human lineage.

It is also possible that the down-regulation of many ASD-linked genes is less deleterious than the down-regulation of a single gene. As an analogy, whole-genome duplications can be well-tolerated in vertebrates, even though duplication of some individual genes—including many of those linked to ASD—can be far more deleterious. An intuitive explanation for this counter-intuitive observation is that relative expression levels, or stoichiometry, could impact fitness even more than absolute expression levels[91]. Under this model, the key idea is that the down-

503 regulation of many ASD-linked genes would have less impact on their relative levels than a
504 change in the expression of a single gene. Excitingly, CRISPR-based methods to precisely
505 manipulate the expression levels of many genes at once may soon allow us to more directly test
506 this hypothesis. Overall, it will be important to develop a deeper understanding of how cell types
507 and genes implicated in ASD and SCZ have evolved in the human lineage as this will improve
508 our understanding of uniquely human traits and neuropsychiatric disorders.

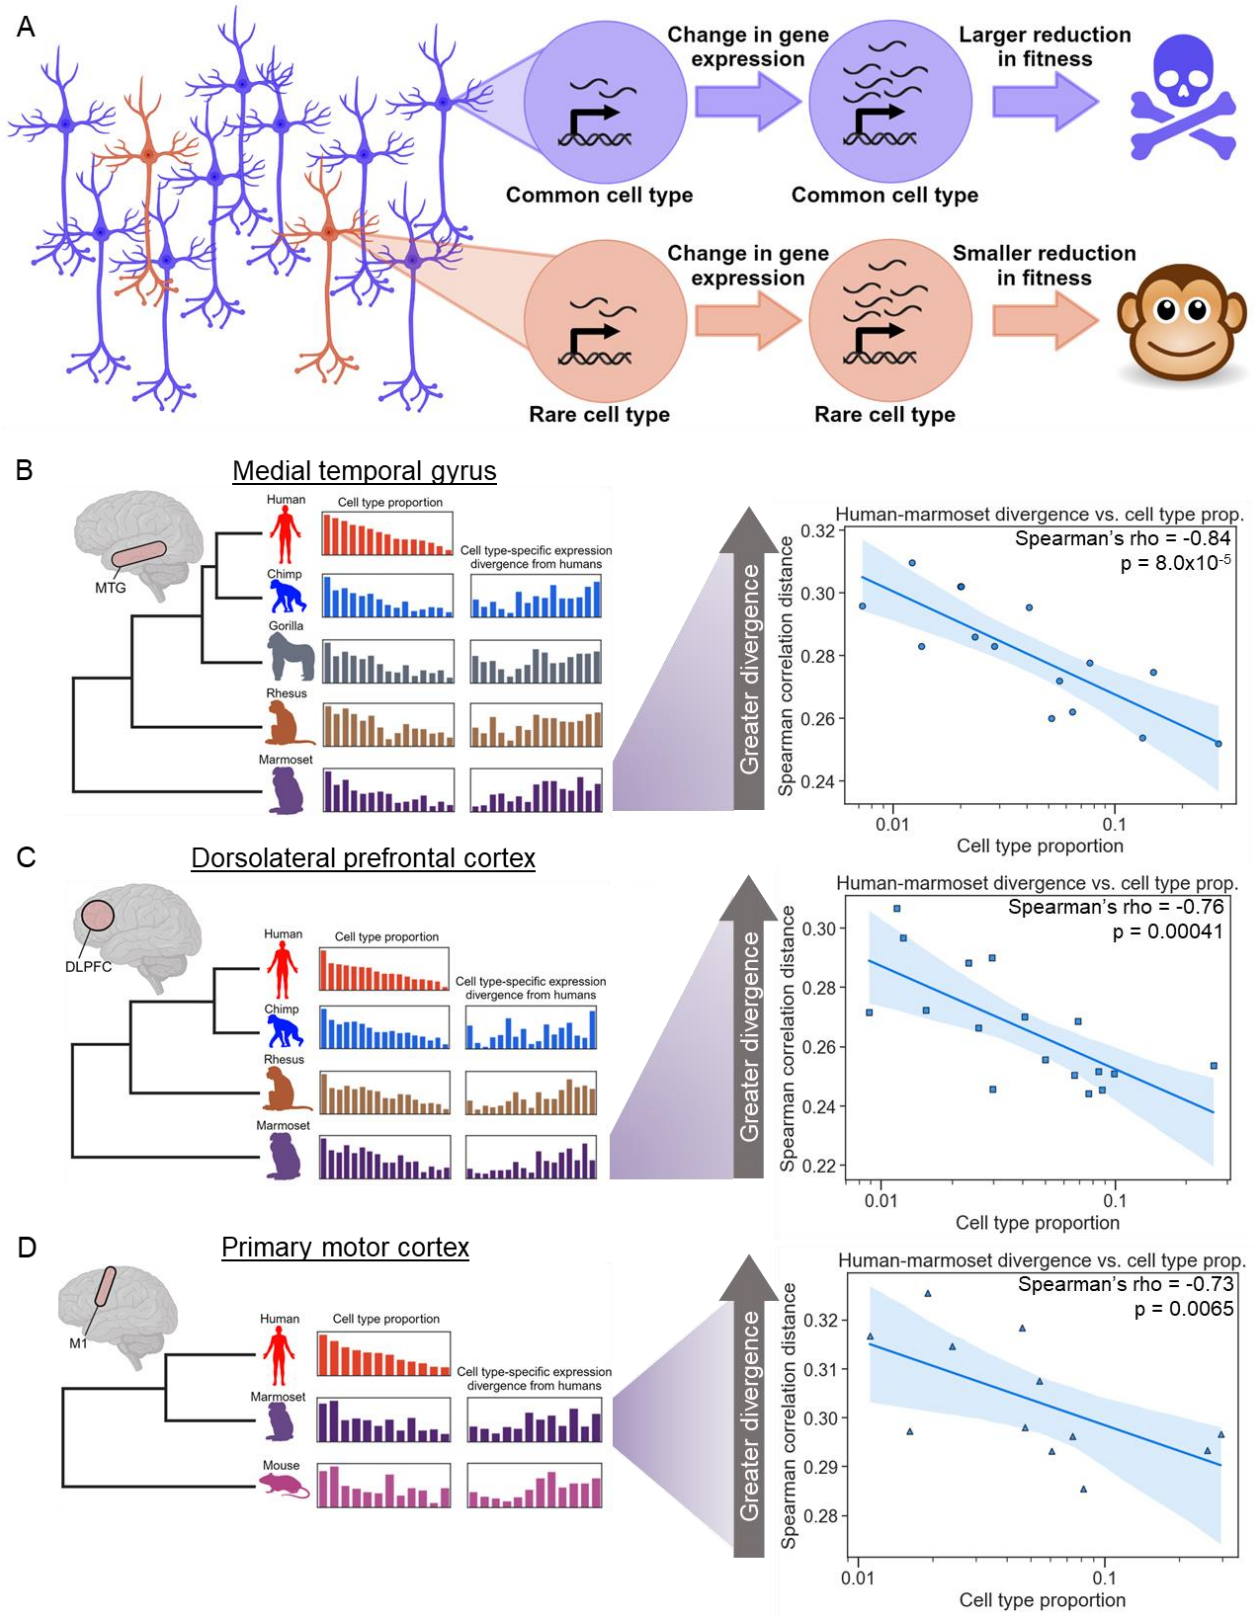

Figure 1: More common neuronal cell types evolve more slowly than rarer types. A) Rationale for hypothesis that more common neuronal types might evolve more slowly than rarer types. A gene expression change in a common cell type has a large negative effect on fitness whereas the same change in a rarer cell type has a smaller effect. **B) On the left:** outline of data analysis strategy. SnRNA-seq from the MTG of five species (14 subclasses of neuron) was analyzed and used to measure cell type proportion and pairwise divergence between species. **On the right:** plot showing the correlation between neuronal subclass proportion ($\log_{10}$ scale on the x-axis) and subclass-specific divergence between human and marmoset in the MTG. A representative iteration from 100 independent down-samplings is shown. The Spearman's rho and p-value shown are the median across 100 independent down-samplings (see methods for details). The line and shaded region are the line of best fit from a linear regression and 95% confidence interval respectively. **C)** Same as (B) but snRNA-seq from the DLPFC (17 subclasses of neuron) of four species was analyzed. **D)** Same as (B) but snRNA-seq from M1 (12 subclasses of neuron) of three species was analyzed.

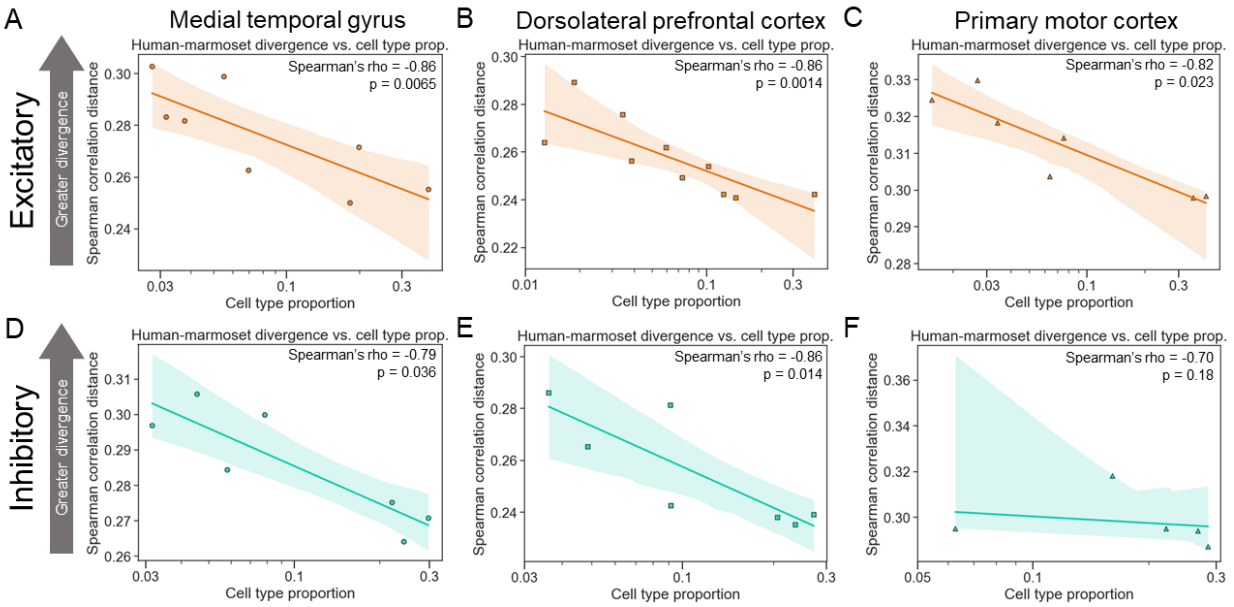

Figure 2: More common neuronal cell types evolve more slowly than rarer types within excitatory and inhibitory classes. A) Plot showing the correlation between neuronal subclass proportion (log_{10} scale on the x-axis) and subclass-specific divergence between human and marmoset in the MTG, restricted to excitatory neurons. A representative iteration from 100 independent down-samplings is shown. The Spearman's rho and p-value shown are the median across 100 independent down-samplings (see methods for details). The line and shaded region are the line of best fit from a linear regression and 95% confidence interval respectively. **B)** Same as in (A) but for the DLPFC data. **C)** Same as in (A) but for the M1 data. **D)** Same as in (A) but restricting to inhibitory neurons. **E)** Same as in (B) but restricting to inhibitory neurons. **F)** Same as in (C) but restricting to inhibitory neurons.

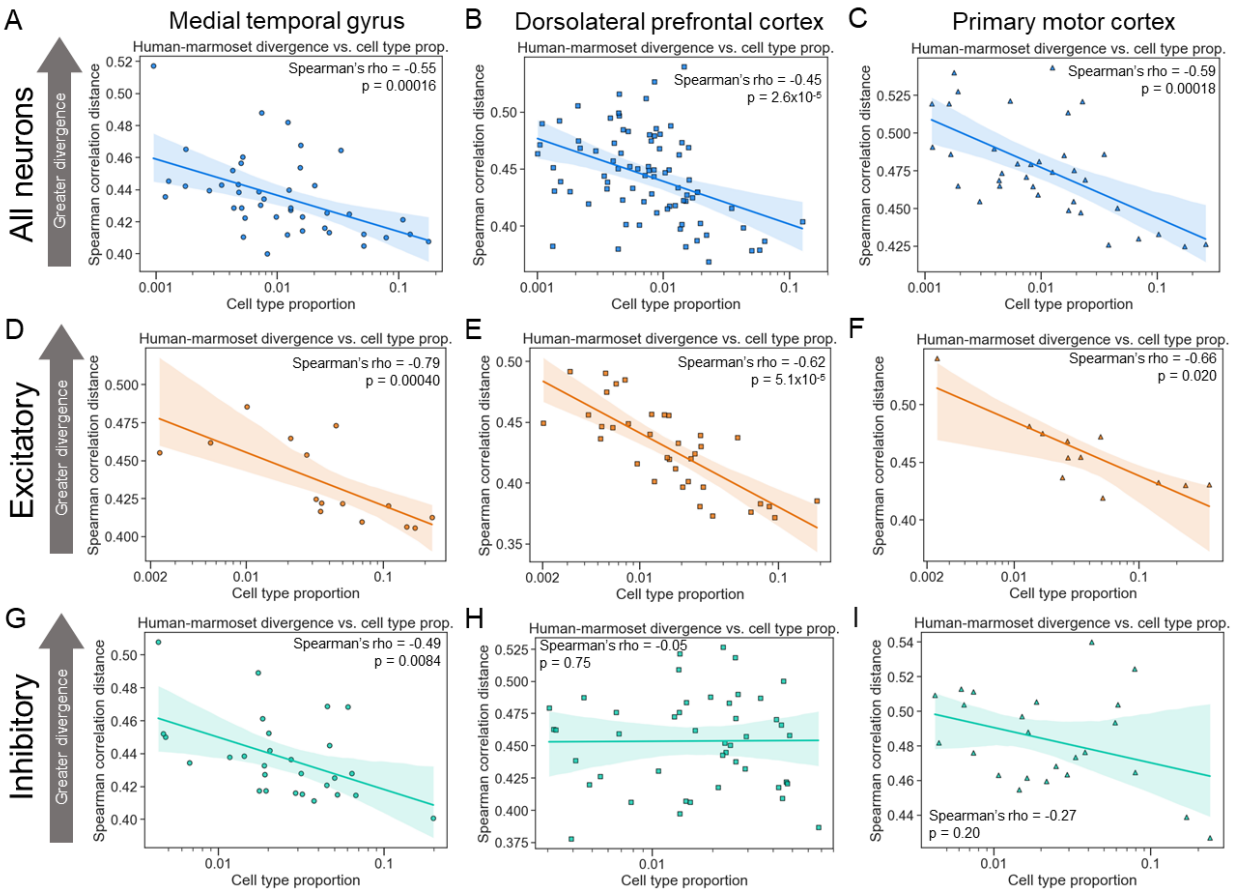

Figure 3: More common neuronal cell types evolve more slowly than rarer types at the subtype level. A) Plot showing the correlation between neuronal subtype proportion (log_{10} scale on the x-axis) and subtype-specific divergence between human and marmoset in the MTG. A representative iteration from 100 independent down-samplings is shown. The Spearman's rho and p-value shown are the median across 100 independent down-samplings (see methods for details). The line and shaded region are the line of best fit from a linear regression and 95% confidence interval respectively. **B)** Same as in (A) but for the DLPFC data. **C)** Same as in (A) but for the M1 data. **D)** Same as in (A) but restricting to excitatory neurons. **E)** Same as in (B) but restricting to excitatory neurons. **F)** Same as in (C) but restricting to excitatory neurons. **G)** Same as in (A) but restricting to inhibitory neurons. **H)** Same as in (B) but restricting to inhibitory neurons. **I)** Same as in (C) but restricting to inhibitory neurons.

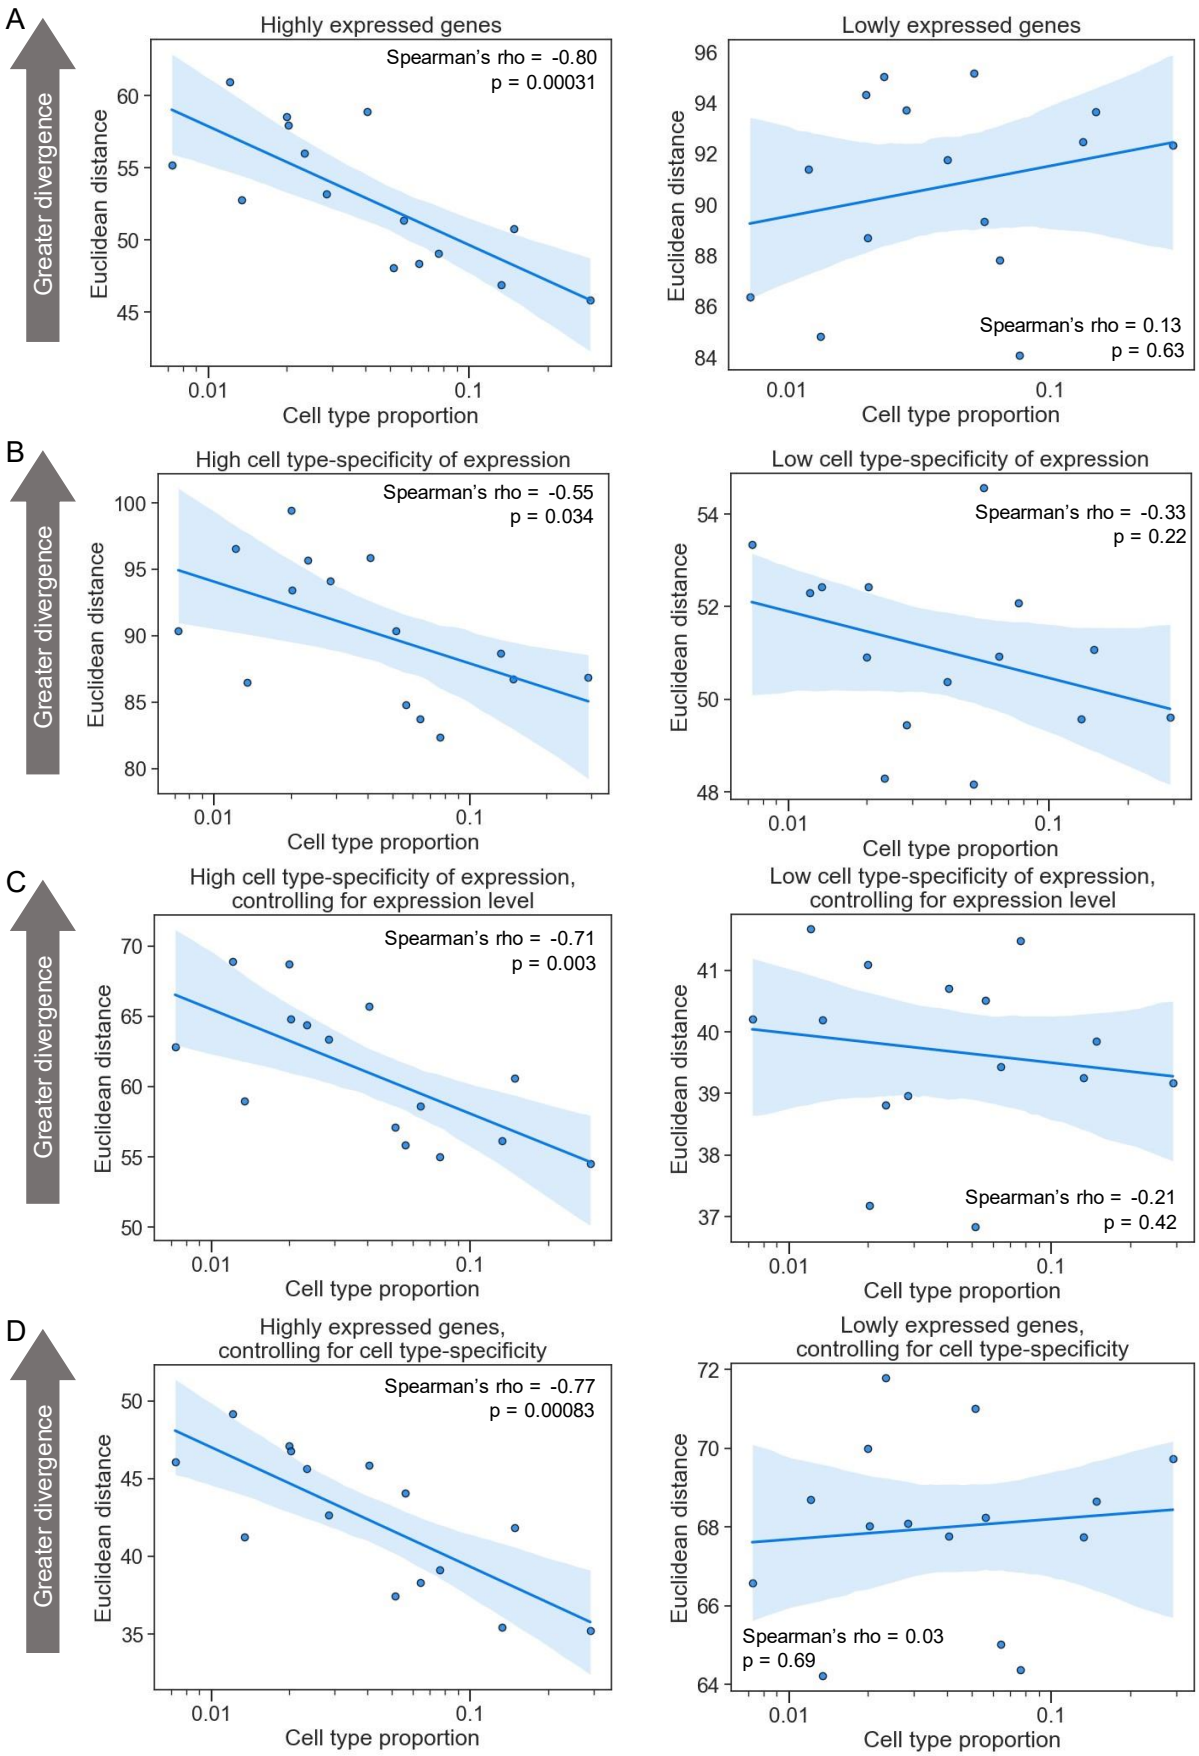

Figure 4: More highly expressed, cell type-specific genes drive the negative correlation between cell type proportion and evolutionary divergence. A) Left: Plot showing the correlation between neuronal subtype proportion (log_{10} scale on the x-axis) and subtype-specific divergence for highly expressed genes between human and marmoset in the MTG. A representative iteration from 100 independent down-samplings is shown. The Spearman's rho and p-value shown are the median across 100 independent down-samplings (see methods for details). The line and shaded region are the line of best fit from a linear regression and 95% confidence interval respectively. **Right:** Same as the left but for lowly expressed genes. **B) Left:** Same as in (A) but for genes with more cell type-specific expression. **Right:** Same as left but for genes with less cell type-specific expression. **C)** Same as in (A) but controlling for expression level (Methods). **D)** Same as in (B) but controlling for cell type-specificity of expression.

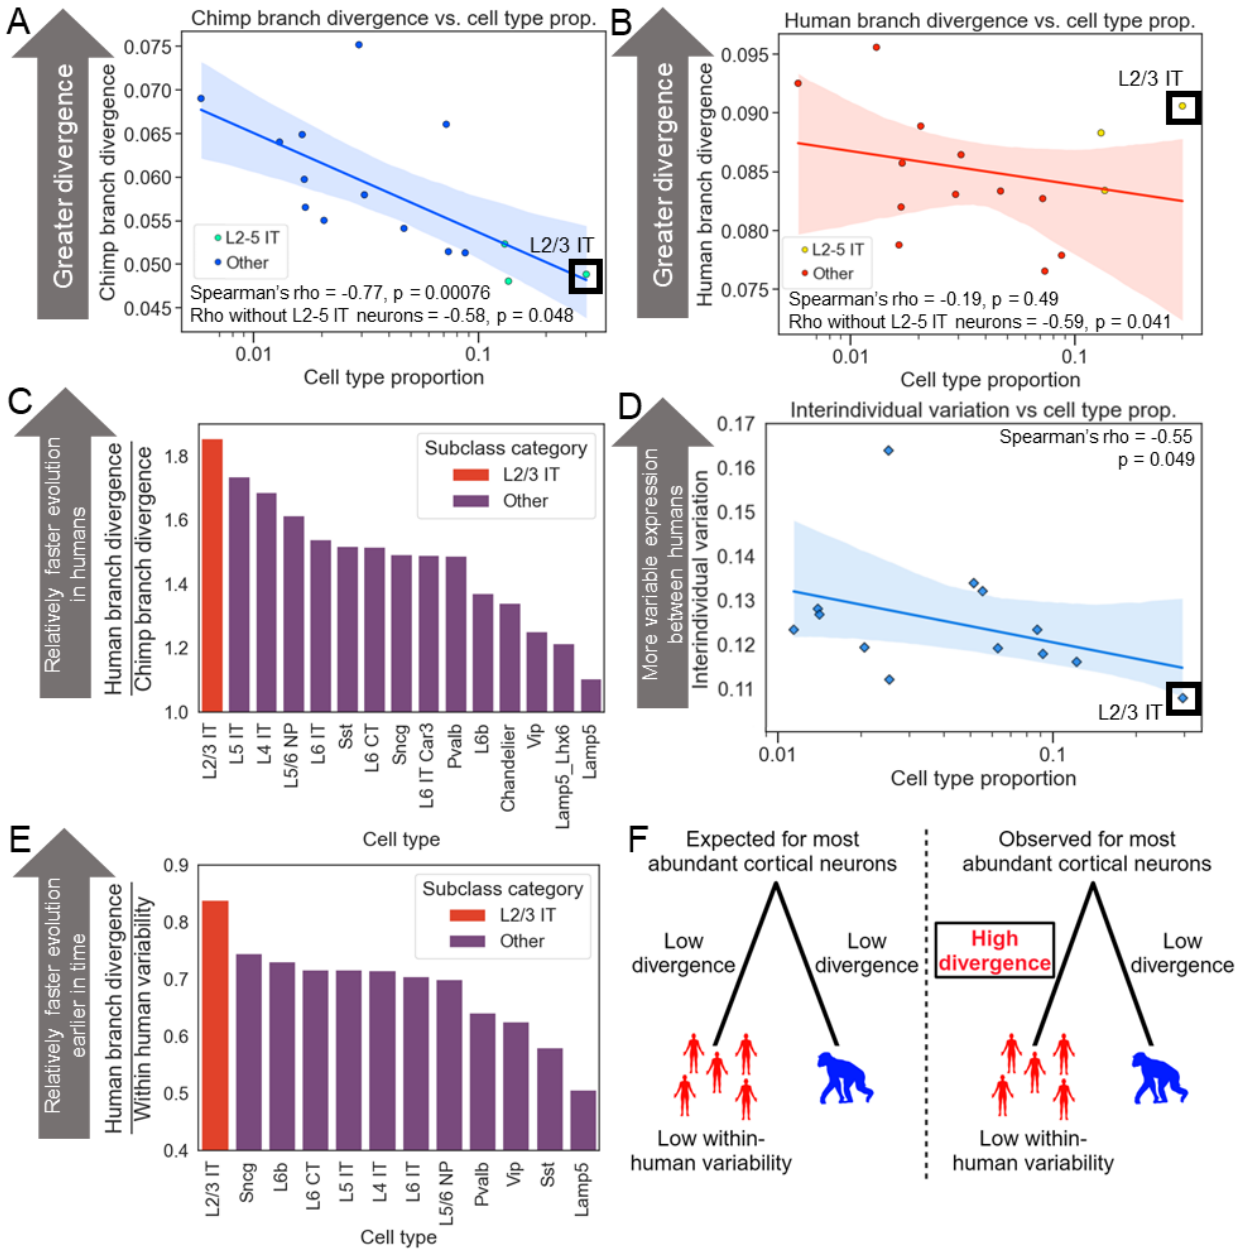

Figure 5: Accelerated evolution of L2/3 IT neurons in the human lineage. A) Plot showing the correlation between neuronal subclass proportion (log_{10} scale on the x-axis) and subclass-specific divergence on the chimp branch in the MTG. Chimp branch divergence was computed for each of 100 down-samplings and the mean across those down-samplings is shown. The line and shaded region are the line of best fit from a linear regression and 95% confidence interval respectively. Green points indicate L2-5 IT neurons. **B)** Same as in (A) but for human branch

569 divergence. Yellow points indicate L2-5 IT neurons. **C)** Barplot showing the human branch
570 divergence divided by the chimp branch divergence for each subclass. **D)** Plot showing the
571 correlation between neuronal subclass proportion (log_{10} scale on the x-axis) and subclass-
572 specific interindividual variation across DLPFC samples from 25 human individuals. A
573 representative iteration from 100 independent down-samplings is shown. The Spearman's rho
574 and p-value shown are the median across 100 independent down-samplings (see methods for
575 details). The line and shaded region are the line of best fit from a linear regression and 95%
576 confidence interval respectively. **E)** Barplot showing the human branch divergence divided by
577 the within human variability for each subclass. **F)** Conceptual model for accelerated evolution of
578 L2/3 IT neurons in the human lineage.
579

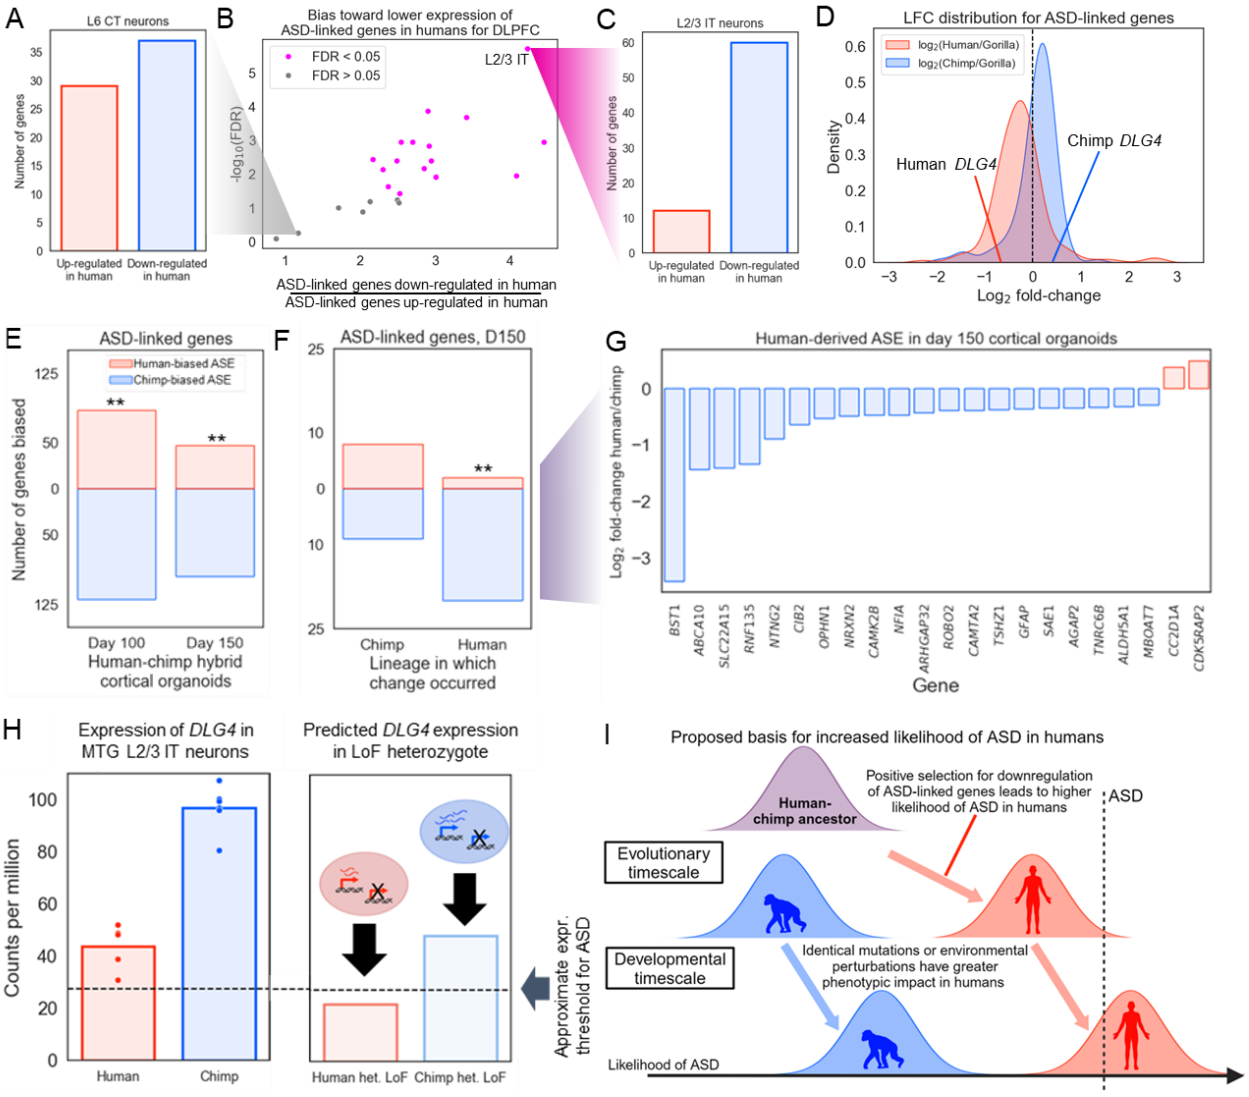

Figure 6: Positive selection for down-regulation of ASD-linked genes in the human lineage. A) Barplot showing the number of high confidence ASD-linked genes that are up-regulated in human and number of genes that are down-regulated in human relative to chimp in DLPFC L6 CT neurons. **B)** Volcano plot showing the fold-enrichment for down-regulation in human DLPFC (x-axis) and the -$\log_{10}$ binomial FDR (y-axis). Subclasses with FDR < 0.05 are shown in red; only subclasses with at least 500 differentially expressed (DE) genes up-regulated in human and 500 differentially expressed genes down-regulated in human are shown. **C)** Barplot showing the number of high confidence ASD-linked genes that are up-regulated in

589 human and number of genes that are down-regulated in human relative to chimp in DLPFC L2/3
590 IT neurons. **D)** Distribution of $\log_2$ fold-changes (x-axis) comparing human or chimpanzee to
591 gorilla in L2/3 IT neurons for high confidence ASD-linked genes with FDR < 0.05 when
592 comparing human and chimpanzee. Only genes with absolute $\log_2$ fold-change less than 3 are
593 shown. **E)** Barplot showing the number of differentially expressed ASD-linked genes with higher
594 allele-specific expression from the human allele (red) and higher expression from the chimp
595 allele (blue) in cortical organoids. ** indicates binomial p < 0.01. **F)** Barplot showing the number
596 of differentially expressed ASD-linked genes with higher allele-specific expression from the
597 human allele (red) and higher expression from the chimp allele (blue) in day 150 cortical
598 organoids for human-derived and chimp-derived genes separately. ** indicates binomial p <
599 0.01. **G)** Plot showing the $\log_2$ allele-specific expression ratios of differentially expressed,
600 human-derived, ASD-linked genes in day 150 cortical organoids. **H)** Left: Expression of *DLG4* in
601 MTG L2/3 IT neurons. Right: Predicted expression of *DLG4* if one copy of the gene were non-
602 functional. **I)** Conceptual model for how positive selection for down-regulation of ASD linked
603 genes led to higher likelihood of ASD in humans compared to chimpanzees.
604

Methods

Quantifying cell type-specific gene expression divergence between species

We analyzed three main datasets in this study, which we refer to by the cortical area sampled (MTG, DLPFC, M1). These were the only studies meeting both of our inclusion criteria: multiple species profiled in the same study using the same snRNA-seq protocols for each species within a study, and at least 10 orthologous cell types having 250 or more cells per species. As an example of a study that did not meet these inclusion criteria, we can consider a recent multi-species atlas of the retina[92]. While this has a sufficient number of cells, species, and orthologous cell types, different protocols were used for different species and not all species were sampled as part of the same study. For example, different antibodies were used to enrich for subpopulations of cells in different species and some species did not have a sufficient number of cells profiled without enrichment to accurately estimate cell type proportions.

All statistical tests and analyses were performed in python using scipy v1.10.1[93] except for the DESeq2 analysis. For the M1 and MTG data, we converted from RDS files to h5 files using Seurat and Seurat Disk[94]. We conducted all analyses within each dataset to avoid batch effects from comparing across datasets. We used the cell type annotations and counts matrices directly from the study that first reported the dataset in conjunction with scanpy v1.7.2[95]. The procedure outlined below was performed 100 times independently on each dataset unless otherwise noted. To quantify cell type-specific expression divergence without confounding with cell type proportion, we first down-sampled the number of cells in each cell type so that it was equal across all cell types and species. We down-sampled without replacement to 250 cells at the subclass level and 50 cells at the subtype level for the main analysis presented in the text. Only

subclasses and subtypes with at least this many cells were included in downstream analysis. We then restricted to 5-way one-to-one protein-coding non-mitochondrial orthologs (downloaded from ensembl biomart for hg38)[96] between human, chimpanzee, gorilla, rhesus macaque, and marmoset for the MTG and DLPFC data and 3-way one-to-one orthologs for human, marmoset, and mouse for the M1 dataset. We then summed expression across all cells within a cell type to create a pseudobulked expression profile for that cell type.

For each possible pairwise comparison between species, we down-sampled the total counts in each cell type so that it was equal across all cell types for both species in the comparison. We then computed counts per million (CPM) in each cell type. After computing CPM, we filtered out genes with (1) fewer than 25 counts in both species or (2) fewer than 1 CPM in both species per cell type. As a result, if a gene passed the filtering criteria in one cell type but not another it would be included only for the cell type in which it passed the filtering criteria. We then computed the $\log_2$(CPM) and used the Spearman correlation distance to measure the gene expression divergence between species in each cell type.

Notably, this process involved several analysis decisions that could affect our results. To test how robust our results were to these choices, we tested all combinations of the following:

1. Down-sampling to 50, 100, 250, or 500 cells.
2. Filtering genes with fewer than 5, 10, 25, or 50 counts.
3. Filtering genes with fewer than 1 or 5 CPM.
4. Using $\log_2$(CPM) or not log transforming.
5. Using the Spearman correlation distance, Pearson correlation distance, Euclidean distance, or L1 distance metrics.

In general, our results were robust to any combination of these parameters (Supplemental Tables). When stratifying, we only used a subset of these combinations due to the greater number of computations required.

Computing cell type proportions and correlation with gene expression divergence

All three datasets were generated with single-nucleus RNA-sequencing (snRNA-seq) and so likely accurately represent the true proportion of neuronal cell types in the neocortex[97]. To compute cell type proportions, we restricted to neuronal cells with greater than or equal to the number of cells we down-sampled to. We then computed cell type proportion separately for each species by dividing the number of cells of each type by the total number of cells profiled. For each interspecies comparison, we averaged the cell type proportion across both species. We then computed the Spearman correlation between the averaged cell type proportions and cell type-specific gene expression divergence computed as described above. As we did this across 100 independent down-samplings (numbered 1 to 100), we reported the median Spearman's rho and p-value throughout the text and figures. If there was an individual down-sampling iteration that had the median Spearman's rho and p-value, we made the scatterplots shown in Figures 1-4 using the first such iteration. If no iteration had the median rho and p-value, we showed the iteration closest to the median with the greatest number of iterations that had that rho and p-value. For example, if 22 iterations resulted in rho = -0.5 and 19 iterations resulted in rho = -0.6, both of which were closest to the median of -0.55, then an iteration with -0.5 would be shown. If there was still a tie after this process, we showed the iteration with the lowest number. Because the Spearman correlation is a nonparametric rank-based test, it is unaffected by any rank-preserving transformation of the data; therefore our choice to show

scatter plots with log-transformed cell type proportions was for visualization only and had no effect on the results.

To estimate divergence along the human branch, we used the formula:

$$\frac{HC\ divergence + HG\ divergence - CG\ divergence}{2}$$

Here, HC stands for human-chimp, HG stands for human-gorilla, and CG stands for chimp-gorilla.

Similarly, to estimate divergence along the chimp branch, we used the formula:

$$\frac{HC\ divergence + CG\ divergence - HG\ divergence}{2}$$

Stratifying by expression level, cell type-specificity of expression, and constraint on expression

To stratify by expression level, we ranked genes by the average CPM between the two species being compared for each cell type separately. We then assigned the top third of genes with the highest expression to the highly expressed bin, the next third to the moderately expressed bin, and the remaining third to the lowly expressed bin. Whenever we stratified by expression level or another metric, we used the Euclidean distance to measure gene expression divergence because the limited dynamic range of expression for the moderately and lowly expressed bins led to unrealistically high correlation distances. Similarly, we ranked genes by Tau[62], a measure of how cell type-specifically a gene is expressed, and split those genes into three bins. We computed Tau separately for both species across all subclasses or subtypes with a sufficient number of cells and then computed the average value for each gene. For constraint on

expression, we considered all genes with heterozygous fitness effect[61] s_{het} > 0.1 to be highly constrained, genes with s_{het} between 0.1 and 0.01 as moderately constrained, and the remaining genes with s_{het} < 0.01 to be lowly constrained. Because there was a different number of genes in each bin in this case, we down-sampled genes to reach an equal number in each bin.

When controlling for expression level and stratifying by Tau, we compared the high bin with the moderate and low bins separately. To control for expression, we first computed the $\log_2$ fold-change between all genes in the high bin and all genes in the moderate or low bin and restricted to pairs of genes with absolute $\log_2$ fold-change less than 0.05. We then split this list of gene pairs into those with a negative $\log_2$ fold-change, positive $\log_2$ fold-change, and zero $\log_2$ fold-change, shuffled the list, and removed duplicate genes. We kept all gene pairs with a $\log_2$ fold-change of zero and down-sampled the list of gene pairs with positive or negative $\log_2$ fold-change so that there were an equal number in each category. This resulted in a final set of genes in the high bin with matched expression to genes in the moderate or low bin which we used to compute cell type-specific gene expression divergence. When controlling for Tau, we applied the same strategy but required an absolute $\log_2$ fold-change less than 0.01.

Comparing interindividual variability in gene expression and cell type proportion

To measure the within-human interindividual variation in cell type-specific gene expression, we used a uniformly processed dataset from the DLPFC[63]. We restricted to control samples from individuals of European ancestry with an age of death greater than or equal to 25. We selected thirteen neuronal subclasses for which the majority of individuals had greater than 50 nuclei profiled for further analysis and restricted to samples with greater than or equal to 50 nuclei for

all thirteen subclasses. After this filtering process, 25 samples remained. Next, we down-sampled to 50 nuclei from each subclass in each dataset and computed pseudobulked counts. We then down-sampled counts so that there was an equal number of total counts across all subclasses for each individual. For each subclass, we removed genes with average counts across all individuals less than 25 and computed CPM. We then computed the Spearman correlation distance between each sample and the mean expression profile across all samples and took the mean of those 25 correlation distances as our measure of cell type-specific gene expression variation within humans. We performed this procedure across 100 independent down-samplings. To estimate cell type proportions, we computed the cell type proportions for the thirteen subclasses and averaged them together. We then computed the Spearman correlation between the subclass-specific interindividual variation and the cell type proportions across the 100 down-samplings. We report the median Spearman's rho and p-value across the 100 down-samplings and show the first down-sampling with the median Spearman's rho and p-value in Figure 5D.

Analysis of ASD- and SCZ-linked genes in snRNA-seq data

We used the SFARI gene database of ASD-linked genes and considered any genes with a score of 1 to be "high-confidence" (233 total) and all genes regardless of score to be all ASD-linked genes (1176 genes)[98]. As we are not aware of a similar resource for SCZ, we used the 31 genes with FDR < 0.1 in a recent rare variant association study for SCZ[76]. Throughout, FDRs were corrected for multiple tests with the Benajmini-Hochberg method. To identify differentially expressed (DE) genes and compute $\log_2$ fold-changes between species, we ran DESeq2[99] on the subclass-level pseudobulked counts and used apeglm[100] to shrink the $\log_2$ fold-changes. To test for a bias toward lower expression of ASD- and SCZ-linked genes in each cell type, we

restricted to genes with FDR < 0.05 in the human-chimpanzee comparison and use the binomial test comparing the number of genes with negative $\log_2$ fold-change (i.e. higher expression in chimpanzee) to the number of genes with positive $\log_2$ fold-change. We used the frequency of negative $\log_2$ fold-changes among all genes with FDR < 0.05 as the background probability in the binomial test. We repeated this for both high-confidence and all ASD-linked genes.

To determine whether the higher expression in chimpanzees relative to human was more likely due to changes on the chimpanzee branch or the human branch, we first filtered to only high-confidence ASD-linked genes that were differentially expressed between chimpanzees and gorillas in L2/3 IT neurons. Genes were assigned as having a significant human-derived or chimpanzee-derived expression change in the MTG dataset by comparison with the human-gorilla and chimpanzee-gorilla $\log_2$ fold-changes. First, if the absolute human-gorilla and chimpanzee-gorilla $\log_2$ fold-change both were greater than the absolute human-chimpanzee $\log_2$ fold-change, that gene was considered ambiguous. After removing ambiguous genes, a gene was considered as having a human-derived expression change if the absolute human-gorilla $\log_2$ fold-change was greater than the absolute human-chimpanzee $\log_2$ fold-change and vice versa for chimpanzee-derived.

Analysis of ASD-linked genes in human-chimpanzee hybrid cortical organoid data

We used the previously described dataset from human-chimpanzee cortical organoids, reprocessed as previously described[86]. Briefly, reads were aligned to the human (hg38) and chimpanzee (PanTro6) genomes with STAR and corrected for mapping bias using Hornet[101]. Reads were assigned to the human or chimpanzee allele using a set of high-confidence human-chimp single nucleotide differences and collapsed to counts per gene with ASEr. DESeq2[99] was

779 used to identify genes with significant ASE with the hybrid line that each sample was from used
780 as a covariate. DESeq2[99] and apeglm[100] were used to compute $\log_2$ fold-changes. For the below
781 analyses, we used the chimpanzee-aligned data, which has a very slight bias toward higher
782 expression from the human allele, to ensure that our analyses were conservative.
783
784 To test for a significant bias toward down or up-regulation from the human allele for ASD- or
785 SCZ-linked genes, we restricted to genes with FDR < 0.05 in the cortical organoid data and
786 intersected those genes with the list of ASD- or SCZ-linked genes. We then used the binomial
787 test comparing the number of genes with negative $\log_2$ fold-change (i.e. higher expression in
788 chimpanzee) to the number of genes with positive $\log_2$ fold-change. We used the frequency of
789 negative $\log_2$ fold-changes among all genes with FDR < 0.05 as the background probability in
790 the binomial test. We repeated this for both high-confidence and all ASD-linked genes. To
791 investigate whether these *cis*-regulatory changes likely occurred in the human or chimpanzee
792 lineage, we used the assignments as human- or chimpanzee-derived from L2/3 IT neurons in
793 the MTG dataset described above. For genes that had matching human-chimpanzee $\log_2$ fold-
794 change sign in both the MTG and cortical organoid datasets, we created a 2x2 table of
795 human/chimp-derived and down/up-regulated from the human allele and applied Fisher's exact
796 test.
797

798 ***Analysis of constraint on ASD-linked genes in humans and chimpanzees***

799

800 We used previously published dN/dS computations[70] and restricted only to genes with at least
801 one synonymous and nonsynonymous difference on both the human and chimpanzee
802 branches. We compared the dN/dS for ASD-linked genes with a paired t-test. To compute the
803 number of genetic differences within 5 kilobases of the transcription start site (TSS) for each

804 lineage, we used our previously described set of high-confidence human-chimpanzee single
805 nucleotide genetic differences[86]. Briefly, this was created by identifying all single nucleotide
806 differences between PanTro6 and hg38 and the filtering out sites that were not homozygous for
807 the reference allele in 3 humans and 3 chimpanzees. We then intersected this with a previously
808 described list of human-chimpanzee orthologous TSS expanded by 2.5 kilobases on either side
809 and restricted to only TSS for ASD-linked genes[87]. To correct for the slightly larger number of
810 human-derived sites across all genes, we down-sampled the human-derived variants near the
811 TSS of ASD-linked genes, keeping a fraction of sites equal to the total number of chimp-derived
812 genetic differences divided by the total number of human-derived genetic differences. We then
813 used a paired t-test to compare the two distributions.
814
815 To compare the within-species variance for humans and chimpanzees in expression of ASD-
816 linked genes, we computed the variance in pseudobulked CPM from L2/3 IT neurons across
817 individuals in the DLPFC and MTG separately. As the mean expression level and batch effects
818 can have a major impact on expression variance, we normalized the variance to the variance of
819 the 100 genes with closest mean expression to each ASD-linked gene. To do this, we computed
820 the fraction of those 100 genes with smaller variance than the focal ASD-linked gene in each
821 species and dataset separately. We then compared the values in human and chimpanzee with a
822 paired t-test.
823
824 ***Analysis of postsynaptic proteomics data***
825
826 We plotted PSD-95 protein abundances from the supplemental materials of Wang et al[74]. We
827 used the t-test to compare levels between species.
828

Acknowledgements: We thank Liqun Luo and other Luo lab members for helpful discussion. We also thank Leslie Magtanong and other members of the Fraser Lab for helpful discussions and feedback on the manuscript. Some subfigures were made with biorender.

Funding: Funding was provided by NIH R01HG012285 (awarded to HBF). ALS was supported by a fellowship under grant number FA9550-21-F-0003.

Authors contributions: ALS performed all bioinformatic analysis, visualization, validation, and writing of software with guidance from HBF. ALS and HBF wrote the main and ALS created the figures with input from HBF. HBF provided funding for the study.

Competing interests: All authors declare no competing interests.

Data availability: The MTG data is available from https://labshare.cshl.edu/shares/gillislab/resource/Primate_MTG_coexp/Great_Ape_Data/. The metadata for the MTG study is available from https://github.com/AllenInstitute/Great_Ape_MTG/blob/master/data/ (files ending in "for_plots_and_sharing_12_16_21.RDS"). The DLPFC data is available from https://data.nemoarchive.org/biccn/grant/u01_sestan/sestan/transcriptome/sncell/10x_v3/. The M1 data is available from https://data.nemoarchive.org/publication_release/Lein_2020_M1_study_analysis/Transcriptomics/sncell/10X/. The constraint metric s_{het} was downloaded from the supplemental materials of https://www.biorxiv.org/content/10.1101/2023.05.19.541520v1. The constraint metric pLI was downloaded from https://gnomad.broadinstitute.org/downloads#v4-constraint. The SFARI ASD-linked genes were downloaded from https://gene.sfari.org/. The SCZ-linked genes were downloaded from https://www.nature.com/articles/s41586-022-04556-w. Protein abundance

855 measurements in the post-synaptic density of humans, rhesus macaques, and mice were
856 obtained from the supplemental materials of https://www.nature.com/articles/s41586-023-
857 06542-2. The human population DLPFC single nucleus RNA-seq used to compute within human
858 cell type-specific gene expression variation was downloaded from
859 https://brainscope.gersteinlab.org/output-sample-annotated-matrix.html. The data from the study
860 of human-chimpanzee hybrid cortical organoids was downloaded from
861 https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE144825. dN/dS for the human and
862 chimp lineages was downloaded from https://doi.org/10.1186/1471-2164-15-599. All code
863 needed to reproduce the analyses described in this study is available at
864 https://github.com/astarr97/Cell_Type_Evolution.
865
866 **Materials and correspondence:**
867 Correspondence should be addressed to HBF at hbfraser@stanford.edu and ALS at
868 astarr97@stanford.edu. Requests for materials will be fulfilled by the corresponding authors.
869
870 **List of supplementary materials:**
871
872 Supplemental Figures 1-41
873 Supplemental Tables 1-4
874 Supplemental Table 1 contains the median Spearman's rho and p-value for the correlation
875 between cell type divergence and proportion across a variety of parameter combinations.
876 Supplemental Table 2 contains the median Spearman's rho and p-value for the correlation
877 between cell type divergence and proportion stratifying by expression level across a variety of
878 parameter combinations.

879 Supplemental Table 3 contains the median Spearman's rho and p-value for the correlation
880 between cell type divergence and proportion stratifying by evolutionary constraint across a
881 variety of parameter combinations.
882 Supplemental Table 4 contains the median Spearman's rho and p-value for the correlation
883 between cell type divergence and proportion stratifying by cell type-specificity of expression
884 across a variety of parameter combinations.
885

References

1. Zeisel, A. *et al.* Brain structure. Cell types in the mouse cortex and hippocampus revealed by single-cell RNA-seq. *Science* **347**, 1138–1142 (2015).
2. Tasic, B. *et al.* Adult mouse cortical cell taxonomy revealed by single cell transcriptomics. *Nat. Neurosci.* **19**, 335–346 (2016).
3. Yao, Z. *et al.* A high-resolution transcriptomic and spatial atlas of cell types in the whole mouse brain. *Nature* **624**, 317–332 (2023).
4. Krienen, F. M. *et al.* Innovations present in the primate interneuron repertoire. *Nature* **586**, 262–269 (2020).
5. Jorstad, N. L. *et al.* Comparative transcriptomics reveals human-specific cortical features. *Science* **382**, eade9516 (2023).
6. Hodge, R. D. *et al.* Conserved cell types with divergent features in human versus mouse cortex. *Nature* **573**, 61–68 (2019).
7. Bakken, T. E. *et al.* Comparative cellular analysis of motor cortex in human, marmoset and mouse. *Nature* **598**, 111–119 (2021).
8. Ma, S. *et al.* Molecular and cellular evolution of the primate dorsolateral prefrontal cortex. *Science* **377**, eabo7257 (2022).
9. Eyre-Walker, A. Evolutionary genomics. *Trends Ecol. Evol.* **14**, 176 (1999).
10. Pál, C., Papp, B. & Hurst, L. D. Highly expressed genes in yeast evolve slowly. *Genetics* **158**, 927–931 (2001).
11. Hirsh, A. E. & Fraser, H. B. Protein dispensability and rate of evolution. *Nature* **411**, 1046–1049 (2001).
12. Fraser, H. B., Hirsh, A. E., Steinmetz, L. M., Scharfe, C. & Feldman, M. W. Evolutionary rate in the protein interaction network. *Science* **296**, 750–752 (2002).

13. Duret, L. & Mouchiroud, D. Determinants of substitution rates in mammalian genes: expression pattern affects selection intensity but not mutation rate. *Mol. Biol. Evol.* **17**, 68–74 (2000).
14. Drummond, D. A. & Wilke, C. O. Mistranslation-induced protein misfolding as a dominant constraint on coding-sequence evolution. *Cell* **134**, 341–352 (2008).
15. Drummond, D. A., Raval, A. & Wilke, C. O. A single determinant dominates the rate of yeast protein evolution. *Mol. Biol. Evol.* **23**, 327–337 (2006).
16. Drummond, D. A., Bloom, J. D., Adami, C., Wilke, C. O. & Arnold, F. H. Why highly expressed proteins evolve slowly. *Proc. Natl. Acad. Sci. U. S. A.* **102**, 14338–14343 (2005).
17. Yang, Z. PAML 4: Phylogenetic Analysis by Maximum Likelihood. *Mol. Biol. Evol.* **24**, 1586–1591 (2007).
18. Arendt, D. *et al.* The origin and evolution of cell types. *Nat. Rev. Genet.* **17**, 744–757 (2016).
19. Pembroke, W. G., Hartl, C. L. & Geschwind, D. H. Evolutionary conservation and divergence of the human brain transcriptome. *Genome Biol.* **22**, 52 (2021).
20. Kebschull, J. M. *et al.* Cerebellar nuclei evolved by repeatedly duplicating a conserved cell-type set. *Science* **370**, eabd5059 (2020).
21. Tosches, M. A. *et al.* Evolution of pallium, hippocampus, and cortical cell types revealed by single-cell transcriptomics in reptiles. *Science* **360**, 881–888 (2018).
22. Peng, Y.-R. *et al.* Molecular Classification and Comparative Taxonomics of Foveal and Peripheral Cells in Primate Retina. *Cell* **176**, 1222-1237.e22 (2019).
23. Luo, L. Architectures of neuronal circuits. *Science* **373**, eabg7285 (2021).
24. Jagadeesh, K. A. *et al.* Identifying disease-critical cell types and cellular processes by integrating single-cell RNA-sequencing and human genetics. *Nat. Genet.* **54**, 1479–1492 (2022).
25. Wightman, D. P. *et al.* A genome-wide association study with 1,126,563 individuals identifies new risk loci for Alzheimer's disease. *Nat. Genet.* **53**, 1276–1282 (2021).

26. Jansen, I. E. *et al.* Genome-wide meta-analysis identifies new loci and functional pathways influencing Alzheimer's disease risk. *Nat. Genet.* **51**, 404–413 (2019).

27. Galakhova, A. A. *et al.* Evolution of cortical neurons supporting human cognition. *Trends Cogn. Sci.* **26**, 909–922 (2022).

28. Berg, J. *et al.* Human neocortical expansion involves glutamatergic neuron diversification. *Nature* **598**, 151–158 (2021).

29. Kanton, S. *et al.* Organoid single-cell genomic atlas uncovers human-specific features of brain development. *Nature* **574**, 418–422 (2019).

30. Dear, R. *et al.* Cortical gene expression architecture links healthy neurodevelopment to the imaging, transcriptomics and genetics of autism and schizophrenia. *Nat. Neurosci.* **27**, 1075–1086 (2024).

31. Parikshak, N. N. *et al.* Integrative functional genomic analyses implicate specific molecular pathways and circuits in autism. *Cell* **155**, 1008–1021 (2013).

32. Wamsley, B. *et al.* Molecular cascades and cell type–specific signatures in ASD revealed by single-cell genomics. *Science* **384**, eadh2602 (2024).

33. Velmeshev, D. *et al.* Single-cell genomics identifies cell type-specific molecular changes in autism. *Science* **364**, 685–689 (2019).

34. Pintacuda, G. *et al.* Protein interaction studies in human induced neurons indicate convergent biology underlying autism spectrum disorders. *Cell Genomics* **3**, 100250 (2023).

35. Batiuk, M. Y. *et al.* Upper cortical layer–driven network impairment in schizophrenia. *Sci. Adv.* **8**, eabn8367 (2022).

36. Trubetskoy, V. *et al.* Mapping genomic loci implicates genes and synaptic biology in schizophrenia. *Nature* **604**, 502–508 (2022).

957 37. Ruzicka, W. B. *et al.* Single-cell multi-cohort dissection of the schizophrenia transcriptome.
958 *Science* **384**, eadg5136 (2024).
959 38. Sullivan, P. F., Yao, S. & Hjerling-Leffler, J. Schizophrenia genomics: genetic complexity and
960 functional insights. *Nat. Rev. Neurosci.* (2024) doi:10.1038/s41583-024-00837-7.
961 39. Jutla, A., Foss-Feig, J. & Veenstra-VanderWeele, J. Autism spectrum disorder and schizophrenia:
962 An updated conceptual review. *Autism Res. Off. J. Int. Soc. Autism Res.* **15**, 384–412 (2022).
963 40. Dodell-Feder, D., Tully, L. M. & Hooker, C. I. Social impairment in schizophrenia: new
964 approaches for treating a persistent problem. *Curr. Opin. Psychiatry* **28**, 236–242 (2015).
965 41. Sato, M., Nakai, N., Fujima, S., Choe, K. Y. & Takumi, T. Social circuits and their dysfunction in
966 autism spectrum disorder. *Mol. Psychiatry* **28**, 3194–3206 (2023).
967 42. Lugo Marín, J. *et al.* Prevalence of Schizophrenia Spectrum Disorders in Average-IQ Adults with
968 Autism Spectrum Disorders: A Meta-analysis. *J. Autism Dev. Disord.* **48**, 239–250 (2018).
969 43. Lai, M.-C. *et al.* Prevalence of co-occurring mental health diagnoses in the autism population: a
970 systematic review and meta-analysis. *Lancet Psychiatry* **6**, 819–829 (2019).
971 44. Zheng, S. *et al.* Autistic traits in first-episode psychosis: Rates and association with 1-year
972 recovery outcomes. *Early Interv. Psychiatry* **15**, 849–855 (2021).
973 45. Sikela, J. M. & Searles Quick, V. B. Genomic trade-offs: are autism and schizophrenia the steep
974 price of the human brain? *Hum. Genet.* **137**, 1–13 (2018).
975 46. Crow, T. J. Is schizophrenia the price that Homo sapiens pays for language? *Schizophr. Res.* **28**,
976 127–141 (1997).
977 47. Zug, R. & Uller, T. Evolution and dysfunction of human cognitive and social traits: A
978 transcriptional regulation perspective. *Evol. Hum. Sci.* **4**, e43 (2022).
979 48. Yoshida, K. *et al.* Single-neuron and genetic correlates of autistic behavior in macaque. *Sci.*
980 *Adv.* **2**, e1600558 (2016).

49. Faughn, C. *et al.* Brief Report: Chimpanzee Social Responsiveness Scale (CSRS) Detects Individual Variation in Social Responsiveness for Captive Chimpanzees. *J. Autism Dev. Disord.* **45**, 1483–1488 (2015).

50. Marrus, N. *et al.* Initial description of a quantitative, cross-species (chimpanzee-human) social responsiveness measure. *J. Am. Acad. Child Adolesc. Psychiatry* **50**, 508–518 (2011).

51. MacLean, E. L. Unraveling the evolution of uniquely human cognition. *Proc. Natl. Acad. Sci. U. S. A.* **113**, 6348–6354 (2016).

52. Mody, M., MGH/HST Athinoula A. Martinos Center for Biomedical Imaging, Harvard Medical School, Department of Radiology, Charlestown, MA, Belliveau, J. W., & MGH/HST Athinoula A. Martinos Center for Biomedical Imaging, Harvard Medical School, Department of Radiology, Charlestown, MA. Speech and Language Impairments in Autism: Insights from Behavior and Neuroimaging. *Am. Chin. J. Med. Sci.* **5**, 157 (2012).

53. Chang, X. *et al.* Language abnormalities in schizophrenia: binding core symptoms through contemporary empirical evidence. *Schizophrenia* **8**, 95 (2022).

54. Doan, R. N. *et al.* Mutations in Human Accelerated Regions Disrupt Cognition and Social Behavior. *Cell* **167**, 341-354.e12 (2016).

55. Shin, T. *et al. Rare Variation in Noncoding Regions with Evolutionary Signatures Contributes to Autism Spectrum Disorder Risk*. http://medrxiv.org/lookup/doi/10.1101/2023.09.19.23295780 (2023) doi:10.1101/2023.09.19.23295780.

56. Pollard, K. S. *et al.* Forces Shaping the Fastest Evolving Regions in the Human Genome. *PLoS Genet.* **2**, e168 (2006).

57. Liu, X. *et al.* Disruption of an Evolutionarily Novel Synaptic Expression Pattern in Autism. *PLOS Biol.* **14**, e1002558 (2016).

1004 58. van den Heuvel, M. P. *et al*. Evolutionary modifications in human brain connectivity associated
1005 with schizophrenia. *Brain J. Neurol.* **142**, 3991–4002 (2019).
1006 59. Burns, J. K. An evolutionary theory of schizophrenia: cortical connectivity, metarepresentation,
1007 and the social brain. *Behav. Brain Sci.* **27**, 831–855; discussion 855-885 (2004).
1008 60. Ploeger, A. & Galis, F. Evolutionary approaches to autism- an overview and integration. *McGill J.*
1009 *Med. MJM Int. Forum Adv. Med. Sci. Stud.* **13**, 38 (2011).
1010 61. Zeng, T., Spence, J. P., Mostafavi, H. & Pritchard, J. K. Bayesian estimation of gene constraint
1011 from an evolutionary model with gene features. *Nat. Genet.* (2024) doi:10.1038/s41588-024-
1012 01820-9.
1013 62. Yanai, I. *et al*. Genome-wide midrange transcription profiles reveal expression level
1014 relationships in human tissue specification. *Bioinforma. Oxf. Engl.* **21**, 650–659 (2005).
1015 63. Emani, P. S. *et al*. Single-cell genomics and regulatory networks for 388 human brains. *Science*
1016 **384**, eadi5199 (2024).
1017 64. Agoglia, R. M. *et al*. Primate cell fusion disentangles gene regulatory divergence in
1018 neurodevelopment. *Nature* **592**, 421–427 (2021).
1019 65. Gokhman, D. *et al*. Human-chimpanzee fused cells reveal cis-regulatory divergence underlying
1020 skeletal evolution. *Nat. Genet.* **53**, 467–476 (2021).
1021 66. Orr, H. A. Testing natural selection vs. genetic drift in phenotypic evolution using quantitative
1022 trait locus data. *Genetics* **149**, 2099–2104 (1998).
1023 67. Fraser, H. B. Genome-wide approaches to the study of adaptive gene expression evolution:
1024 Systematic studies of evolutionary adaptations involving gene expression will allow many
1025 fundamental questions in evolutionary biology to be addressed. *BioEssays* **33**, 469–477 (2011).
1026 68. Wang, B., Starr, A. L. & Fraser, H. B. *Cell Type-Specific* Cis *-Regulatory Divergence in Gene*
1027 *Expression and Chromatin Accessibility Revealed by Human-Chimpanzee Hybrid Cells*.

1028 http://biorxiv.org/lookup/doi/10.1101/2023.05.22.541747 (2023)

1029 doi:10.1101/2023.05.22.541747.

1030 69. Simon, N. M., Kim, Y., Bautista, D. M., Dutton, J. R. & Brem, R. B. Stem cell transcriptional

1031 profiles from mouse subspecies reveal *cis* -regulatory evolution at translation genes. Preprint at

1032 https://doi.org/10.1101/2023.07.18.549406 (2023).

1033 70. Gayà-Vidal, M. & Albà, M. Uncovering adaptive evolution in the human lineage. *BMC Genomics*

1034 **15**, 599 (2014).

1035 71. Chen, S. *et al.* A genomic mutational constraint map using variation in 76,156 human genomes.

1036 *Nature* **625**, 92–100 (2024).

1037 72. Satterstrom, F. K. *et al.* Large-Scale Exome Sequencing Study Implicates Both Developmental

1038 and Functional Changes in the Neurobiology of Autism. *Cell* **180**, 568-584.e23 (2020).

1039 73. Rodríguez-Palmero, A. *et al.* DLG4-related synaptopathy: a new rare brain disorder. *Genet. Med.*

1040 **23**, 888–899 (2021).

1041 74. Wang, L. *et al.* A cross-species proteomic map reveals neoteny of human synapse

1042 development. *Nature* **622**, 112–119 (2023).

1043 75. Autism Spectrum Disorder Working Group of the Psychiatric Genomics Consortium *et al.*

1044 Identification of common genetic risk variants for autism spectrum disorder. *Nat. Genet.* **51**,

1045 431–444 (2019).

1046 76. Singh, T. *et al.* Rare coding variants in ten genes confer substantial risk for schizophrenia.

1047 *Nature* **604**, 509–516 (2022).

1048 77. Lim, L., Mi, D., Llorca, A. & Marín, O. Development and Functional Diversification of Cortical

1049 Interneurons. *Neuron* **100**, 294–313 (2018).

1050 78. Molyneaux, B. J., Arlotta, P., Menezes, J. R. L. & Macklis, J. D. Neuronal subtype specification in

1051 the cerebral cortex. *Nat. Rev. Neurosci.* **8**, 427–437 (2007).

79. Vasseur, E. & Quintana-Murci, L. The impact of natural selection on health and disease: uses of the population genetics approach in humans. *Evol. Appl.* **6**, 596–607 (2013).

80. Benton, M. L. *et al.* The influence of evolutionary history on human health and disease. *Nat. Rev. Genet.* **22**, 269–283 (2021).

81. Varki, A. Loss of N-glycolylneuraminic acid in humans: Mechanisms, consequences, and implications for hominid evolution. *Am. J. Phys. Anthropol.* **116**, 54–69 (2001).

82. Sabeti, P. C. *et al.* Positive Natural Selection in the Human Lineage. *Science* **312**, 1614–1620 (2006).

83. Chou, H.-H. *et al.* A mutation in human CMP-sialic acid hydroxylase occurred after the *Homo-Pan* divergence. *Proc. Natl. Acad. Sci.* **95**, 11751–11756 (1998).

84. Dankwa, S. *et al.* Ancient human sialic acid variant restricts an emerging zoonotic malaria parasite. *Nat. Commun.* **7**, 11187 (2016).

85. Enard, D., Messer, P. W. & Petrov, D. A. Genome-wide signals of positive selection in human evolution. *Genome Res.* **24**, 885–895 (2014).

86. Starr, A. L., Gokhman, D. & Fraser, H. B. Accounting for cis-regulatory constraint prioritizes genes likely to affect species-specific traits. *Genome Biol.* **24**, 11 (2023).

87. Wang, B., Starr, A. L. & Fraser, H. B. *Cell Type-Specific* Cis *-Regulatory Divergence in Gene Expression and Chromatin Accessibility Revealed by Human-Chimpanzee Hybrid Cells*. http://biorxiv.org/lookup/doi/10.1101/2023.05.22.541747 (2023) doi:10.1101/2023.05.22.541747.

88. Vogindroukas, I., Stankova, M., Chelas, E.-N. & Proedrou, A. Language and Speech Characteristics in Autism. *Neuropsychiatr. Dis. Treat.* **18**, 2367–2377 (2022).

89. DeFelipe, J., Alonso-Nanclares, L. & Arellano, J. Microstructure of the neocortex: Comparative aspects. *J. Neurocytol.* **31**, 299–316 (2002).

90. Sohal, V. S. & Rubenstein, J. L. R. Excitation-inhibition balance as a framework for investigating mechanisms in neuropsychiatric disorders. *Mol. Psychiatry* **24**, 1248–1257 (2019).

91. Darnell, R. B. The Genetic Control of Stoichiometry Underlying Autism. *Annu. Rev. Neurosci.* **43**, 509–533 (2020).

92. Hahn, J. *et al.* Evolution of neuronal cell classes and types in the vertebrate retina. *Nature* **624**, 415–424 (2023).

93. Virtanen, P. *et al.* SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* **17**, 261–272 (2020).

94. Hao, Y. *et al.* Integrated analysis of multimodal single-cell data. *Cell* **184**, 3573-3587.e29 (2021).

95. Wolf, F. A., Angerer, P. & Theis, F. J. SCANPY: large-scale single-cell gene expression data analysis. *Genome Biol.* **19**, 15 (2018).

96. Yates, A. D. *et al.* Ensembl Genomes 2022: an expanding genome resource for non-vertebrates. *Nucleic Acids Res.* **50**, D996–D1003 (2022).

97. Ding, J. *et al.* Systematic comparison of single-cell and single-nucleus RNA-sequencing methods. *Nat. Biotechnol.* **38**, 737–746 (2020).

98. Abrahams, B. S. *et al.* SFARI Gene 2.0: a community-driven knowledgebase for the autism spectrum disorders (ASDs). *Mol. Autism* **4**, 36 (2013).

99. Love, M. I., Huber, W. & Anders, S. Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. *Genome Biol.* **15**, 550 (2014).

100. Zhu, A., Ibrahim, J. G. & Love, M. I. Heavy-tailed prior distributions for sequence count data: removing the noise and preserving large differences. *Bioinformatics* **35**, 2084–2092 (2019).

101. Van De Geijn, B., McVicker, G., Gilad, Y. & Pritchard, J. K. WASP: allele-specific software for robust molecular quantitative trait locus discovery. *Nat. Methods* **12**, 1061–1063 (2015).

1100