Supplementary Information

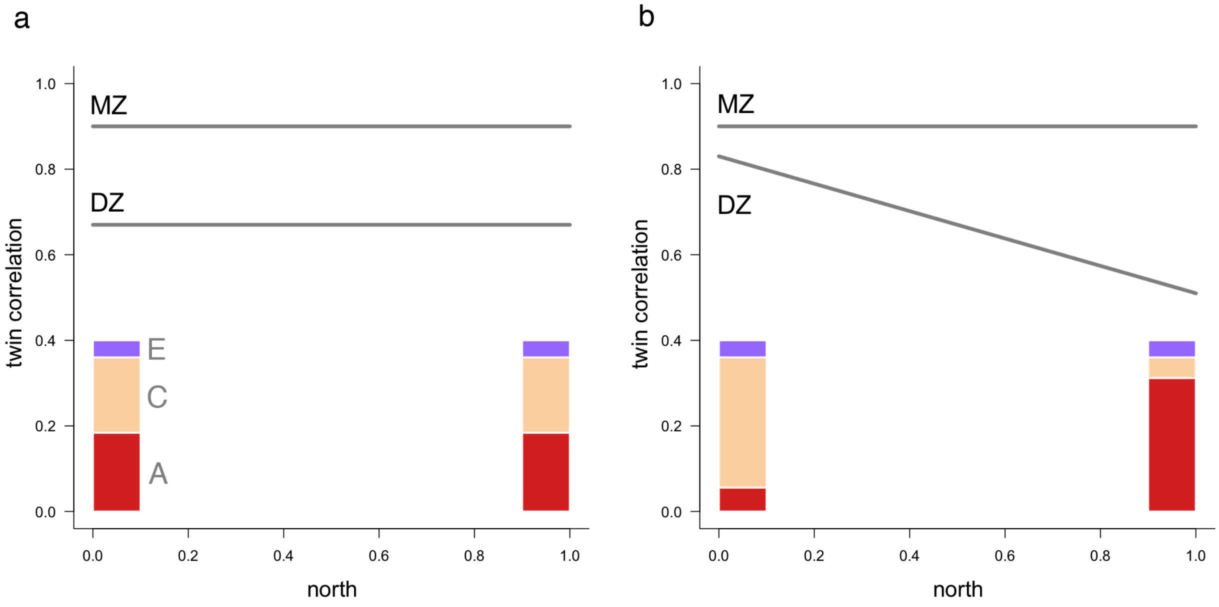

Supplementary Figure 1 | Simulating an artificial map. To validate our approach, we first tested it using simulated data on an artificial square map under two conditions: **a)** variance components are on average constant across the map. To do this, we simulated a random phenotype in twin pairs, keeping the population covariance within MZ (monozygotic) and DZ (dizygotic) twin pairs the same across the map (shown by the labeled lines). There is only chance variation in genetic and environmental components: the bars at the bottom of the graph show that the variance components are the same in the far south (left) and far north (right); in **b)** there is an etiological gradient from south to north across the map. To do this, we again simulated random twin pairs. The population covariance within MZ pairs remains the same across the map. However, while the population covariance for DZ twin pairs is on average the same as in condition a), the simulated correlation within the pairs decreases from south to north. This leads to a systematic increase in the additive genetic (A) component from south to north and a simultaneous decrease in the shared environmental (C) component, indicated by differences in the bars at the bottom of the graph. Because the MZ correlation remains constant, the E component does not change systematically across the map. We applied our geographically sensitive twin model to the simulated data. **Supplementary Figure 2** below shows the results.

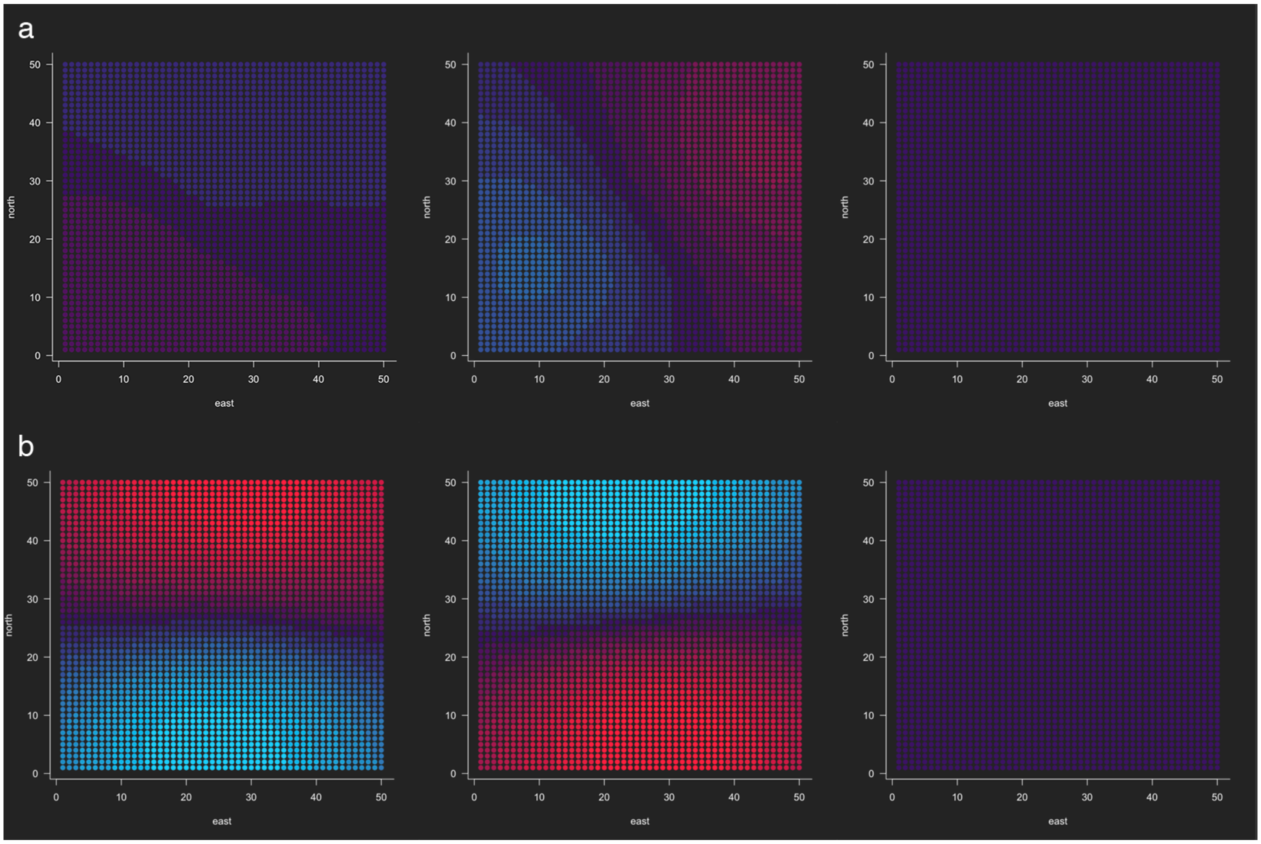

Supplementary Figure 2 | Simulated maps with and without systematic gradients. We plotted the results of the simulation using the same diverging red (high) to blue (low) color gradient we used in the spACE software. Row **a)** represents condition a) from **Supplementary Figure 1**. From left to right, the maps show the additive genetic (A), shared environmental (C) and non-shared environmental (E) variance components. Row **b)** represents condition b) from **Supplementary Figure 1**. We can see some slight chance variation across the maps in the A and C components from condition a). However, these differences are small compared to the intense A and C gradients visible in condition b). As expected, the E map remains a neutral purple in both conditions. These results demonstrate that our geographically sensitive twin model and visualization can reproduce known variation in variance components from simulated data.

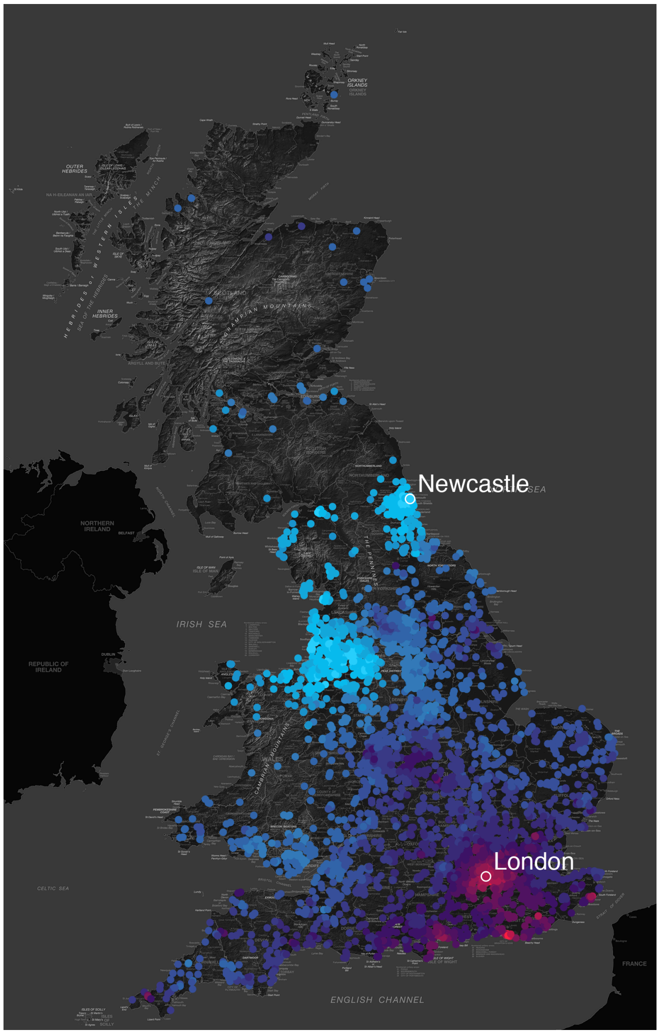

Supplementary Figure 3| Extreme high and low points for the non-shared environmental component of behavior problems. This map reproduces **Figure 3a**, the distribution of environmental influences on behavior problems from blue (low) to red (high). In addition, it picks out two extreme points of the distribution of values: London is an extreme high point in the south-east of England, while Newcastle is an extreme low point in the north-east. These “hot” and “cold” spots are the focus of the simulation described in **Supplementary Figure 4**, below.

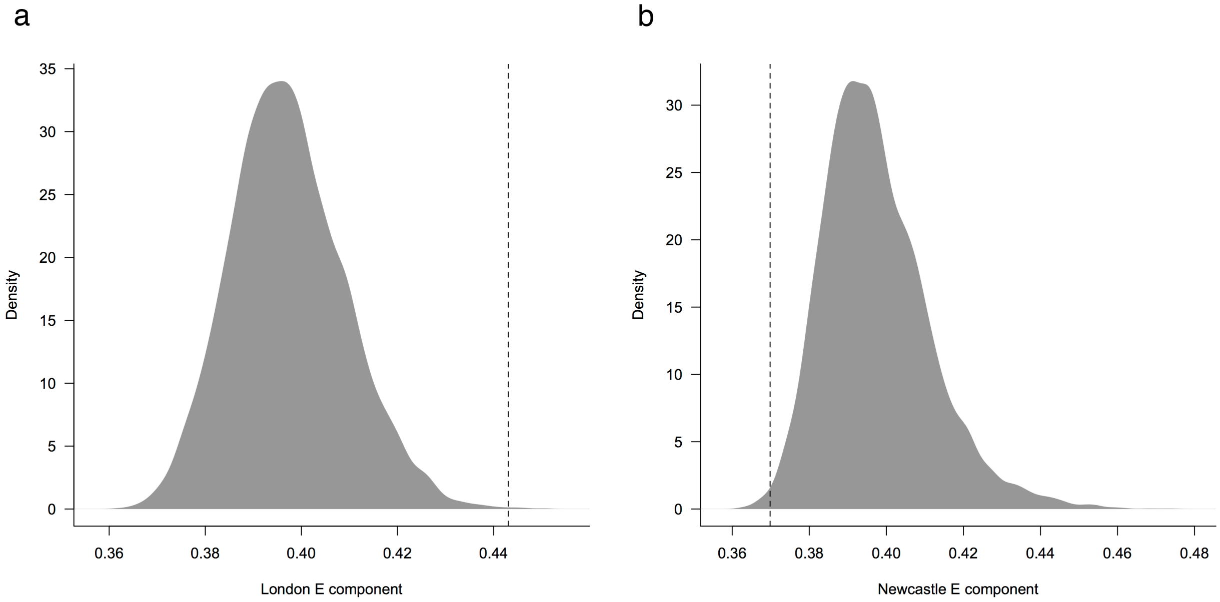

Supplementary Figure 4| Comparing the hotspots observed for behavior problems with simulation under the null. This simulation asked how likely it is that the hot and cold spots identified in **Supplementary Figure 3** occurred by chance. For each point we estimated the value of the non-shared environmental (E) component from the real data. Then we randomized the allocation of twin pairs to geographic locations and ran the geographically sensitive twin model again, estimating the E value from the randomized data; we repeated this process 10,000 times. Graph **a)** shows the distribution of simulated values for London in grey, with the real value represented by the dotted line. Only 7 of the 10,000 simulated values were more extreme than the value estimated from the real data, indicating that this extreme high point is very unlikely to be observed by chance. Likewise, graph **b)** shows the distribution of simulated values for Newcastle in grey, with the real value represented by the dotted line. Only 42 of the 10,000 simulated values were more extreme than the value estimated from the real data, again indicating that this extreme low point is very unlikely to be observed by chance.

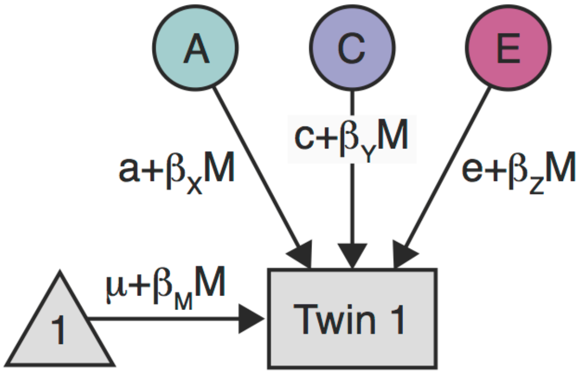

Supplementary Figure 5 | The continuous moderator model. Partial path diagram (showing one twin) for the continuous moderator model described by Purcell (2002; see main text for reference). The model allows additive genetic (*a*), shared environmental (*c*) and non-shared environmental (*e*) paths (and the mean, μ) to vary as a function of a continuous moderator variable (*M*; in our case local variance in household income) that varies from one participant to the next. For example, the contribution of the A variance component is modeled as *a*, plus β times the value of the moderator *M* (which varies). A non-zero β term implies moderation of the variance component. The model estimates both *a* and β_X, making it possible to calculate an estimate of the contribution of additive genetic effects for each value the moderator takes, as shown in **Supplementary Figure 6**.

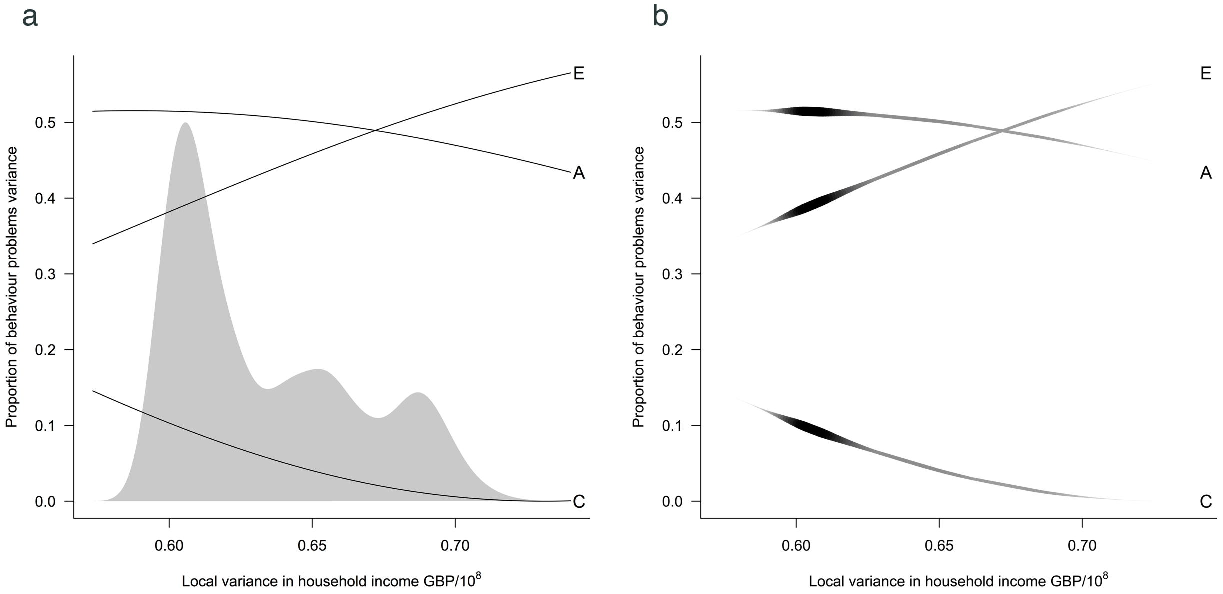

Supplementary Figure 6 | Moderation of behavior problems variance components by local variance in household income. Calculating path coefficients and β estimates by fitting the model described in **Supplementary Figure 5** allows us to estimate genetic and environmental contributions as functions of a measured environment. Graph **a)** shows estimates of additive genetic (A), shared environmental (C) and non-shared environmental (E) influence varying as a function of local variance in household income. The grey density plot behind shows the distribution of the moderator as an indication of the observations supporting each section of the line. Graph **b)** extends a plot first suggested by Purcell (2002; see main text for reference). Rather than varying visibility of the lines as a function of the density of the moderator as Purcell did, we have varied both line width and grey intensity to produce a "quill" plot. The strength of the line clearly indicates information supporting inference at each point, fading out towards the very extremes of the distribution that are supported by fewer observations.

GBP: Great British Pounds.