Supporting Information

Set et al. 10.1073/pnas.1316259111

SI Materials and Methods

Participants. A total of 218 (103 female) undergraduates were recruited from the Behavioral × Biological Economics and Social Sciences (B2ESS) Laboratory at the National University of Singapore. All participants were of ethnic Han Chinese background and had undergone full genome sequencing. A total of 217 (103 female) participants were included in the final analysis after one subject was excluded owing to genotype unavailability.

Procedure. Behavioral data were collected from subjects playing the patent race game in 1-h sessions of 18–24 subjects. In the patent race game, programmed in zTree (1), two players take the role of firms competing to develop a new product. The product is worth a fixed prize and firms are given an endowment to invest. In the asymmetric version of the game we used, the prize is worth \$10 and the two players begin each round with endowments of \$5 and \$4 and are referred to as the "strong" and "weak" players, respectively.

Players can invest any integer amount from their endowment. The investments are subtracted from the potential earnings. To win the prize, one must invest strictly more than the opponent. For example, if the strong player invests \$4 and the weak player invests \$2, the payoff that round to the strong player is \$5 − \$4 + \$10 = \$11, whereas the payoff to the weak player is \$4 − \$2 = \$2. Players' endowments do not carry over from round to round, so the maximum investment available is always either four (for the weak type) or five (for the strong type).

At the beginning of each round, each player was randomly matched with a player of the other type. They played 120 rounds in each role, counterbalanced, for 240 rounds in total. They were fully informed of the rules and matching procedures. Compensation was equal to 10 Singapore dollars (SGD) plus either the average earnings per round or 7 SGD, whichever was higher.

To illustrate how players can anticipate and respond to the actions of others in this game, suppose the weak player observes the strong players frequently investing 5 units. He may subsequently respond by playing 0 to keep his initial endowment. Upon observing this, strong players can exploit the weak player's behavior by investing only 1 unit to obtain both the prize while keeping 4 units from the endowment. This may in turn entice the weak player to move away from investing 0 to win the prize. In contrast, pure reinforcement-learning (RL) players will respond to these changes in opponents' behavior in a much slower manner, because they behave by comparing received payoffs from past investments without consideration for the strategic behavior of others (2).

Genotyping. DNA was extracted from blood samples using QIAamp DNA Blood Midi Kit (Quiagen). SNP genotyping was performed at the Genome Institute of Singapore with HumanOmniExpress-12 v1.0 DNA Analysis Kit (Illumina Inc.). Over 730,000 genetic markers, primarily SNPs, across over 18,000 genes were collected from each subject.

All variable-number tandem repeats (VNTRs) were analyzed with PCR products loaded onto 1.5% (wt/vol) agarose gel with ethidium bromide, run for 1 h at 5 V/cm in Tris/borate/EDTA, and visualized in a UV camera.

The DRD4 exon III VNTR was analyzed with HotStar Plus DNA polymerase (0.3 U per reaction), 1× Q-solution, 1× CoralLoad buffer (Qiagen), 200 μM of each dNTP, 200 nM of each primer, and 10–20 ng of genomic DNA per reaction, in a volume of 10 μL. Primer sequences were as follows: forward 5′- GCGACTACGTGGTCTACTCG -3′, reverse 5′- AGGACCCTCATGGCCTTG -3′ (3). Thermal protocol included an activation step at –95 °C for 5 min, 40 cycles of 94 °C for 30 s, 55 °C for 30 s, 72 °C for 40 s, and final hold at 72 °C for 5 min.

The monoamine oxidase A (MAOA) VNTR was analyzed with PCR ReddyMix Master Mix (Thermo Fisher Scientific), 200 nM of each primer, and 10–20 ng of genomic DNA per reaction, in a volume of 10 μL. Primer sequences were as follows: forward 5′-ACAGCCTGACCGTGGA-3′, reverse 5′- GAACGGACGCTCCATT-3′ (modified from ref. 4). The thermal protocol included an activation step at –95 °C for 5 min, 35 cycles of 95 °C for 30 s, 58 °C for 30 s, 72 °C for 60 s, and final hold at 72 °C for 5 min.

The dopamine transporter (DAT) VNTR was analyzed with PCR ReddyMix Master Mix (Thermo Fisher Scientific), 100 nM of each primer, 0.2% DMSO, and 10–20 ng of genomic DNA per reaction, in a volume of 10 μL. Primer sequences were as follows: forward 5′- TGTGGTGTAGGGAACGGCCTG-3′, reverse 5′- CTTCCTGGAGGTCACGGCTCA-3′ (modified from ref. 5). The thermal protocol included an activation step at –95 °C for 5 min, 35 cycles of 95 °C for 30 s, 61 °C for 30 s, 72 °C for 30 s, and final hold at 72 °C for 10 min.

Gene Selection and Preprocessing. From the dopamine pathway defined in the Kyoto Encyclopedia of Genes and Genomes database, a manually curated collection of pathway maps widely used in gene-set analysis, we included dopamine genes that are involved in (*i*) dopamine synthesis [tyrosine hydroxylase (TH), dopa decarboxylase (DDC), and vesicular monoamine transporter (VMAT)], (*ii*) coding of dopamine receptors (DRD1–5; DRD5 was excluded from the final analysis owing to limited variation of SNPs in the sample), and (*iii*) dopamine transport and clearance [DAT1, catechol-*O*-methyl transferase (COMT), and MAOA/B]. For each gene, SNPs were included according to hg18 coordinates and had minor allele frequency (MAF) exceeding 0.1.

SNP extraction and filtering was conducted using PLINK (6) and snpStats (7). For each gene, SNPs were included if they were contained according to hg18 coordinates and had MAF exceeding 0.1. To reduce dimensionality of the genetic information, we represented each gene as a linear combination of orthogonal vectors using principle component analysis (PCA). Specifically, each analyzed gene is represented by a set of eigenvectors (eigenSNPs) (8) from principal components accounting for at least 90% of the total variation of that gene's SNPs. Occasional genotyping failures (less than 3% of all included SNPs had more than 2 out of 217 failures) were coded with the mean value of the SNP.

X-Chromosome Genes. Because MAOA/B genes reside on the X-chromosome, there is substantial uncertainty regarding the interpretation of allele scores across sex. We addressed this issue in two ways. First, we estimated the model separating sex. Second, we added a sex interaction term to account for multiplicative effects. Both yielded results similar to our original model.

SI Computational Modeling

Base Experience-Weighted Attraction Model (No Genes). Choice behavior was modeled using the hybrid model experience-weighted attraction (EWA) that has been widely used to characterize strategic learning (9). Denote s_i^k as strategy k for player i. Because strategies in the patent race are investments from either a \$5 or \$4 endowment, $k \in \{0, \ldots, 5\}$ when player i is strong and $k \in \{0, \ldots, 4\}$ when player i is weak. For period $t \in \{1, \ldots, 120\}$,

$s_i(t)$ is the amount invested by player i at period t, and $s_{-i}(t)$ is the chosen investment of the opponent at period t.

Player i's (possibly counterfactual) payoff at period t for some s_i^k, given the opponent's actual strategy $s_{-i}(t)$, is equal to the endowment less s_i^k, plus the \$10 prize if $s_i^k > s_{-i}(t)$. This potential payoff is denoted as $\pi_i(s_i^k, s_{-i}(t))$. Notice that, given $s_{-i}(t)$, this potential payoff differs from player i's realized payoff in period t except for when $s_i^k = s_i(t)$.

Player i's expected reward, $V_i^k(t)$, for playing strategy s_i^k in period t is governed by two parameters and updates according to the following:

$$V_i^k(t) = \begin{cases} \dfrac{N(t-1)\cdot\rho\cdot V_i^k(t-1) + \pi_i\left(s_i^k, s_{-i}(t)\right)}{N(t)}, & \text{if} \quad s_i^k = s_i(t) \\ \dfrac{N(t-1)\cdot\rho\cdot V_i^k(t-1) + \delta\cdot\pi_i\left(s_i^k, s_{-i}(t)\right)}{N(t)}, & \text{if} \quad s_i^k \neq s_i(t) \end{cases}, \qquad \textbf{[S1]}$$

where function $N(t) = \rho\cdot N(t-1) + 1$ captures the depreciation of $V_i^k(t)$. If the player believes his opponent is a fast adaptor, he will have a small ρ that depreciates past values faster. In contrast, δ captures the weight between foregone payoffs and actual payoffs when updating values. This corresponds to one of the key insights of the hybrid model that belief learning is equivalent to a model whereby actions are reinforced by foregone payoffs in addition to received payoffs as in RL models. Thus, δ can be interpreted as a psychological inclination toward belief learning (9). That is, the hybrid model reduces to the RL model when $\delta = 0$ and the belief learning model when $\delta = 1$.

To more concretely illustrate the effect of belief learning on behavior, we contrast an EWA strong player with $\delta_1 > 0$ with an RL strong player with $\delta_2 = 0$. Suppose our strong player i invests \$5 and the opponent invests \$1. Both for EWA and RL the value V_i^5 will update to take into account the realized payoff $\pi_i(5,1) = 10$. Unlike the RL player, however, the EWA player with $\delta_1 > 0$ will also update values associated with other actions, even if they were not chosen. For example, in this case the EWA player takes into account the hypothetical payoff $\pi_1(2,1) = 13$ (\$10 prize + \$3 saved from the endowment) based on the opponent's action. Note that as δ_1 increases the greater the sensitivity to the actions of the opponent, ultimately leading to a higher probability that \$2 will be invested in the next round relative to \$5.

Gene-Weighted Model. To account for gene variation, we allowed δ or ρ to vary according to the set of eigenSNPs or VNTR dummy variables. For example, in the case for the DAT1 gene, there were three eigenSNPs, and thus we replace the δ parameter in Eq. **S1** with the individualized term

$$\delta_i^G = \delta_0 + \delta_{E1}\cdot E_{i1} + \delta_{E2}\cdot E_{i2} + \delta_{E3}\cdot E_{i3},$$

where $\{E_{i1}, E_{i2}, E_{i3}\}$ refers to i's three eigenSNP scores and the associated parameters $\{\delta_{E1}, \delta_{E2}, \delta_{E3}\}$ refer to the coefficients on the eigenSNPs. The same procedure is followed for the ρ parameter. Note that this approach implicitly assumes a linear allele–dose–expression–response relationship. We relax this assumption in later analyses by allowing for SNP–SNP interaction.

Behavioral Data Analysis. To calibrate the models given subjects' behavior in the game, we estimated parameters of each model, including initial condition $N(0)$, using subjects' responses by maximizing the logistic log likelihood of the model predictions. To convert values into choices, we used a logit or softmax function to calculate the probability of player i playing strategy k in the next round, $p_i^k(t+1) = e^{\lambda\cdot V_i^k(t)} / \sum_{l=1}^{L} e^{\lambda\cdot V_i^l(t)}$, where λ is an estimated parameter capturing subjects' sensitivity to difference in expected reward associated with the different actions.

Using choice probabilities calculated from the softmax function, we performed maximum likelihood estimation with a grid search over a large range of values for all free parameters in all estimations, because the likelihood function is not globally concave. We aggregated observations conditional on the roles of the subjects and then fit the choice data by maximizing the log likelihood of the observed choices over rounds t for subject i. That is, $\sum_i \sum_t \log(p_i^{s_i(t)}(t))$. Maximum-likelihood estimation of parameters was performed using the quasi-Newton algorithm implemented in the fminunc function in MATLAB. Approximately 100 random or evenly spaced interior starting values were tried, all of which produced essentially identical estimates.

Individual SNP Analysis. We compare our gene-set methodology to other candidate gene approaches by analyzing a selection of individual SNPs for each of the significant genes. These SNPs were identified by cross-referencing the genetic markers available to us with the tagging SNPs suggested by the International HapMap Project's Generic Genome Browser (10). Appropriate tagging SNPs were determined based on pairwise correlations (11). For reference data, we used Han Chinese in Beijing in Data Rel 27 Phase II+III, Feb 09, on NCBI B36 assembly, dbSNP b126. R^2 and MAF cutoffs were 0.8 and 0.1, respectively.

SNP–SNP Interactions. To account for SNP–SNP interactions, we extended the eigenSNP approach by performing PCA on the set of regressors produced from a third-order interaction of the underlying SNP data. For example, if a gene contained 4 SNPs, we performed PCA on the set of 84 regressors, resulting from 4 original SNPs, an additional 16 second-order interaction terms, and a further additional 64 third-order interaction terms. Using the same procedure as outlined above, we took the set of eigenSNPs that explained at least 90% of the variance and included them in our computational model.

Permutation *P* Values. Under the permutation test null hypothesis, individuals are interchangeable, so label-swapping provides a new dataset sampled under the null hypothesis. In each permutation, therefore, the within-gene correlations are preserved and only the behavior–genotype relation is destroyed (6). For each gene, data were permutated 1,000 times by shuffling the gene–subject pairing. The reported P value is equal to the proportion of tests where model fit of the permuted dataset improved upon those of the original, unpermuted dataset.

Empirical *P* Values. Empirical P values were determined by comparing model fit of the gene within the dopamine pathway to comparison genes across the entire genome but outside of the dopamine pathway. A gene was considered comparable if (*i*) it was represented by the same or similar number of SNPs and (*ii*) these SNPs generated the same number of principal components according to the procedure outlined above. A range of SNPs was allowed in cases where an exact match produced too few genes (Table S1). This typically occurred when there were a large number of SNPs in the gene.

Formal Dissociation Test. To formally compare effect size of prefrontal and striatal dopamine genes on choice behavior, we contrasted, using a bootstrap procedure, the mean eigenSNP coefficients for COMT and MAOB against those for DAT1 and DRD2 (12). Specifically, for each of 1,000 iterations of the bootstrap we created a pseudosample by sampling with replacement behavioral and genetic data from 218 participants, and performed maximum likelihood estimates as described above. The resulting coefficients were standardized to ensure

comparability across eigenSNPs, and the reported P value is equal to the proportion of tests where the mean coefficient for one gene set was greater than that of the other gene set.

SI Results

Predictive Accuracy of EWA Model. To assess the ability of our model to capture choice behavior in the patent race, we compared actual proportion of investment against predicted investment proportion (Fig. S1*A*). This is equivalent to a scatterplot of the empirical and EWA prediction proportions as reported in Table S2. Each point represents an investment strategy, that is, strong investment of 5, separately for strong and weak roles. The predicted investment proportion was computed by averaging the round-by-round predictions of the baseline EWA model, aggregating all players over all 120 rounds. The dashed diagonal line represents perfect agreement between the model predictions and actual play. As evident from how closely each point lies to this line, model prediction and actual play are in good agreement, with a χ^2 test result of $P < 10^{-8}$ and a mean difference of less than 5%.

In addition, we sought to incorporate visualization of game dynamics by separating predictions into 30-round blocks, with blocks in the same sequence connected in a series (Fig. S1*B*). All points lie near the diagonal line, confirming the success of the hybrid model of capturing actual play at the finer temporal resolution. The successful modeling of the relative dynamics is also apparent in the generally diagonal pattern within each sequence of points. Although aggregating over rounds and subjects understates the full range of behavior, these plots make clear that the hybrid learning model performs well overall, including the capturing of movements where static approaches are not able to capture.

Note that we do not report a statistic such as R^2 because of the discrete nature of our dependent variable. This issue, as well as model checking techniques such as the one we report above, has been discussed in depth in both neuroimaging and neurophysiological studies of decision making (13).

Incorporating SNP–SNP Interactions. Owing to the low explanatory power of single SNPs, a frequent proposal is that there exists substantial variation that can be explained by accounting for SNP–SNP interactions (14). Accordingly, we investigated this question using our gene-set approach by conducting PCA on regressors formed using third-order interactions of SNPs within a gene (*SI Materials and Methods*). Using the same 90% cutoff rule, we found that incorporating SNP–SNP interactions improved model fit of genes that were previously significant, in particular COMT and DRD2 (Table S4). Interestingly, we did not find qualitative changes in overall level of significance of dopamine genes after accounting for SNP–SNP interactions.

1. Fischbacher U (2007) z-Tree: Zurich toolbox for ready-made economic experiments. *Exp Econ* 10(2):171–178.
2. Hopkins E (2002) Two competing models of how people learn in games. *Econometrica* 70(6):2141–2166.
3. Lichter JB, et al. (1993) A hypervariable segment in the human dopamine receptor D4 (DRD4) gene. *Hum Mol Genet* 2(6):767–773.
4. Sabol SZ, Hu S, Hamer D (1998) A functional polymorphism in the monoamine oxidase A gene promoter. *Hum Genet* 103(3):273–279.
5. Vandenbergh DJ, et al. (1992) Human dopamine transporter gene (DAT1) maps to chromosome 5p15.3 and displays a VNTR. *Genomics* 14(4):1104–1106.
6. Purcell S, et al. (2007) PLINK: A tool set for whole-genome association and population-based linkage analyses. *Am J Hum Genet* 81(3):559–575.
7. Solé X, Guinó E, Valls J, Iniesta R, Moreno V (2006) SNPStats: A web tool for the analysis of association studies. *Bioinformatics* 22(15):1928–1929.
8. Wang K, Abbott D (2008) A principal components regression approach to multilocus genetic association studies. *Genet Epidemiol* 32(2):108–118.
9. Camerer CF, Ho T (1999) Experience-weighted attraction learning in games: A unifying approach. *Econometrica* 67(4):827–874.
10. Gibbs RA, et al.; International HapMap Consortium (2003) The international HapMap project. *Nature* 426(6968):789–796.
11. de Bakker PI, et al. (2005) Efficiency and power in genetic association studies. *Nat Genet* 37(11):1217–1223.
12. Davison AC (1997) *Bootstrap Methods and Their Application* (Cambridge Univ Press, Cambridge, UK).
13. Sugrue LP, Corrado GS, Newsome WT (2005) Choosing the greater of two goods: Neural currencies for valuation and decision making. *Nat Rev Neurosci* 6(5):363–375.
14. Yacubian J, et al. (2007) Gene–gene interaction associated with neural reward sensitivity. *Proc Natl Acad Sci USA* 104(19):8125–8130.

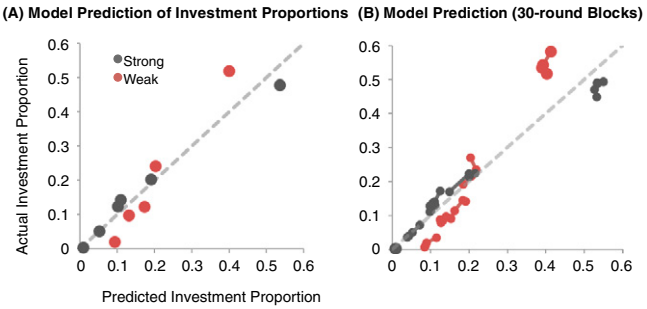

Fig. S1. (*A*) Predicted and actual investment proportions for strong and weak players, averaged over all subjects for all rounds. Each point represents an investment amount (weak, 0–4; strong: 0–5). (*B*) Identical to *A* except separated into 30-round blocks. Blocks are connected by series line.

Table S1. Empirical and predicted choice distributions, calculated as proportion of all players' choices over all rounds

Role	Investment	Nash equilibrium, %	Empirical distribution, %	Conditional win percentage, %	EWA prediction, %
Strong	0	0	0.4	0	0.9
	1	20	20.2	55	19.1
	2	0	5.1	57	5.2
	3	20	12.3	69	10.2
	4	0	14.3	77	11
	5	60	47.8	100	53.6
Weak	0	60	51.9	0	40.0
	1	0	2.0	0	9.4
	2	20	12.2	22	17.3
	3	0	9.7	27	13.1
	4	20	24.2	38	20.3

Table S2. Share of the variance of the gene that can be explained by another gene as calculated using the canonical correlation redundancy index (12)

Gene	COMT,%	DAT,%	DDC,%	DRD1,%	DRD2,%	DRD3,%	DRD4,%	MAOA,%	MAOB,%	TH,%	VMAT2,%
COMT	—	2.6	2.9	0.8	3.1	0.9	0.8	0.2	1.3	0.5	2.8
DAT1	2.1	—	2.3	1.2	2.0	1.9	0.9	0.4	1.6	1.0	3.6
DDC	2.9	2.8	—	1.3	2.5	1.6	0.8	0.8	1.4	1.5	3.2
DRD1	1.0	2.0	1.7	—	2.4	0.6	0.7	0.4	0.9	0.7	2.9
DRD2	2.5	2.0	2.0	1.5	—	1.4	1.1	0.7	1.9	3.0	4.3
DRD3	1.3	3.2	2.1	0.6	2.4	—	1.0	0.1	1.8	0.3	2.4
DRD4	3.2	4.6	3.2	2.0	5.4	3.0	—	0.1	2.1	1.0	1.6
MAOA	0.9	1.9	3.0	1.3	3.7	0.4	0.1	—	10.3	0.7	2.7
MAOB	1.7	2.7	1.8	0.9	3.2	1.8	0.7	3.4	—	1.5	1.5
TH	0.9	2.5	3.0	1.1	7.4	0.4	0.5	0.4	2.3	—	3.0
VMAT2	1.8	3.0	2.1	1.5	3.6	1.2	0.3	0.4	3.4	1.0	—

In the lower diagonal of the matrix, the row variable constitutes the dependent variable, and reversed for the upper diagonal. Note the only gene that explained 10% or more of the variance of another gene was the MAOB gene, which explained 10.3% of MAOA (which resides next to the MAOB gene on the X-chromosome) variation.

Table S3. Selection criteria for comparison genes outside of the dopamine pathway

Gene	SNPs	±
DRD1	5	0
DRD2	17	2
DRD3	6	0
DRD4	1	0
COMT	17	2
DAT1	9	1
MAOA	22	4
MAOB	28	6
TH	2	0
DDC	20	2
VMAT2	16	2

Table S4. Gene properties and estimates in SNP–SNP interactions

				Belief learning (δ)			Discounting (ρ)		
Function	Gene	PCs	% Var	LLR	p_{unc}	p_{perm}	LLR	p_{unc}	p_{perm}
Synthesis	TH	2	0.98	0.95	0.387	0.900	1.03	0.357	0.903
	DDC	8	0.91	21.75	0.000	0.737	28.86	0.000	0.73
	VMAT2	12	0.91	57.59	0.000	0.25	35.34	0.000	0.881
Transport/clearance	DAT1	5	0.91	12.98	0.000	0.646	70.17	0.000	0.019
	COMT	8	0.91	78.32	0.000	0.009	80.29	0.000	0.061
	MAOA	1	0.97	10.60	0.000	0.100	0.08	0.689	0.905
	MAOB	3	0.95	25.56	0.000	0.081	9.87	0.000	0.602
Receptor	DRD1	3	0.99	3.84	0.053	0.819	7.98	0.001	0.72
	DRD2	8	0.91	24.38	0.000	0.63	80.34	0.000	0.07
	DRD3	3	0.97	3.49	0.073	0.841	12.76	0.000	0.50
	DRD4	1	1.00	3.40	0.009	0.358	9.46	0.000	0.19

Owing to the low explanatory power of single SNPs, a frequent proposal is that there exists substantial variation that can be explained by accounting for SNP–SNP interactions. We investigated this question using our gene-set approach by conducting PCA on regressors formed using first-, second-, and third-order interactions of SNPs within a gene. Using a 90% cutoff rule as before, we found that incorporating SNP–SNP interactions improved model fit of a number of genes that were previously significant, particularly COMT and DRD2 (Table S3). However, we did not find that previously insignificant genes became significant after accounting for SNP–SNP interactions. PCA, principal component. % Var, percent of total variance captured by included PCs; LLR, log-likelihood ratio (compared to no-gene baseline); p_{unc}, *P* value using likelihood ratio test; p_{perm}, permutation *P* value (see *SI Materials and Methods*); p_{emp}, empirical *P* value (see *SI Materials and Methods*).