

RESEARCH ARTICLE

Virtual Human Twins for In Silico Clinical Trials in Depression: A Systems Medicine Comparison of Escitalopram, Citalopram and Duloxetine

A patient-level mechanistic comparison using the xISCT platform

Working manuscript prepared for bioRxiv | DEPRESSION

1,154	100	434
unique virtual patients	network solutions	shared response proteins
three molecular cohorts	ensemble-based inference	same direction in all arms

AUTHOR INFORMATION

Jose Manuel Mas^{1,2}, Esther Ramirez², Xavier Daura^{1,3,4,*}

¹ Institut de Biotecnologia i de Biomedicina, Universitat Autònoma de Barcelona, 08193 Bellaterra, Spain.

² Independent Researcher, Barcelona, Spain.

³ CIBER de Bioingeniería, Biomateriales y Nanomedicina (CIBER-BBN), 08041 Barcelona, Spain.

⁴ Institució Catalana de Recerca i Estudis Avançats (ICREA), 08010 Barcelona, Spain.

* Corresponding author: Xavier Daura | Xavier.Daura@uab.cat | ORCID 0000-0001-9235-6730

SCOPE

This manuscript presents a mechanistic in silico study aimed at identifying the proteins and molecular routes underlying simulated response to escitalopram, citalopram and duloxetine. Aggregate response outputs are model-behaviour controls and are not comparative clinical efficacy estimates.

1 Abstract

Major depressive disorder is clinically and molecularly heterogeneous, and similar group-level outcomes may arise through different biological routes. We evaluated xISCT, a virtual-human-twin platform for patient-level in silico clinical trials that combines probabilistic molecular states, a Human Protein Network, an ensemble of semisupervised xModels and individual exposure. A common population was evaluated under escitalopram 10 or 20 mg/day, citalopram 40 mg/day and duloxetine 60 mg/day. The study comprised 1,154 unique virtual patients and 100 retained network solutions from a ~14,700-protein, ~270,000-link network. Retention was 61.0–83.3% across disease cohorts and arms. The fixed molecular-to-clinical transformation achieved full coverage and high prediction correlation ($r=0.97$; RMSE 5.29 percentage points), but did not reproduce the prespecified treatment order ($r=0.47$), and 33 of 102 phenotype-objective checks required review. These limitations preclude comparative-efficacy claims. Mechanistically, both escitalopram doses and citalopram generated the same 254-protein informative topology, while citalopram redistributed sensitivity from SERT (43.90%) towards sigma-1 (28.56%) and other receptors. Duloxetine generated a compact 43-protein topology, with 38 proteins shared across all arms and SERT/NET/DAT sensitivities of 91.46%, 7.17% and 1.37%. VMAT2 and FGF1 were strongly activated in every arm, FGFR1 and CRH separated citalopram from the other treatments, and VAMP2/SCAMP2 were conserved functional sinks. Complete responder tables identified 493–2,014 proteins per arm at $q<0.10$ and a 434-protein concordant four-arm core. Shared candidates included BLNK, serum amyloid P, claudin-5, WNT1, KCNJ8, CD44 and HSPG2; treatment-conditioned baseline candidates included mitochondrial proteins for escitalopram 10 mg and SREBF1/PCSK9/ELK1 for duloxetine, whereas no citalopram baseline protein passed $q<0.10$. Literature concordance supports the biological domains represented by VMAT2, FGF/FGFR1, serum amyloid P/SERT trafficking, claudin-5 and WNT signalling, while the remaining signatures are explicit hypotheses for molecular and prospective clinical validation. xISCT therefore provides a traceable systems-medicine framework for explaining simulated patient heterogeneity, not a substitute for a randomized clinical trial.

Keywords in silico clinical trial; virtual human twin; major depressive disorder; systems medicine; escitalopram; citalopram; duloxetine; mechanism of action; responder heterogeneity; PBPK

2 Introduction

Major depressive disorder (MDD) is a leading cause of disability and a biologically heterogeneous syndrome. The World Health Organization estimated that approximately 322 million people experienced depression in 2023, including 5.2% of adults, with a higher prevalence in women than men (6.2% versus 4.1%). Depression was associated with an estimated 727,000 suicide deaths in 2021, and only about one third of affected people in high-income countries received mental-health treatment (World Health Organization, 2026). These figures coexist with marked variation in symptoms, longitudinal course and treatment response.

Therapeutic heterogeneity is visible in pragmatic clinical studies. In STAR*D, remission decreased across successive treatment steps, showing that a diagnosis and treatment sequence do not define a uniform biological response (Rush et al., 2006). Network meta-analysis establishes average antidepressant efficacy at population level (Cipriani et al., 2018), but randomized controlled trials primarily report endpoints and effect distributions. They rarely explain which patient proteins, pathways or network states generated those outcomes. Blood-based proteomic studies have identified immune, endocrine and metabolic response signals, yet replication and mechanistic interpretation remain difficult (Chan et al., 2016; Gadad et al., 2017).

Virtual patients and virtual human twins offer a complementary scale of explanation. A patient-specific molecular state can receive a treatment-specific exposure and propagate that perturbation through a constrained biological model. Repeating this operation across a population yields an in silico clinical trial in which group behaviour and individual molecular trajectories coexist (Viceconti and Hunter, 2016;

Moingeon et al., 2023). A previous pharmacological ISCT compared lisdexamfetamine and methylphenidate in 2,600 virtual patients by combining PBPK, quantitative systems pharmacology and mechanistic analyses, providing a direct methodological precedent for head-to-head virtual trials (Gutiérrez-Casares et al., 2021).

Mechanistic modelling is already used in regulatory contexts when credibility is matched to a defined context of use. The FDA accepted the University of Virginia/Padova type 1 diabetes simulator in 2008 as a substitute for specified preclinical animal studies of insulin-delivery and closed-loop algorithms; subsequent regulatory documents described investigational-device approvals supported by simulation. The European Medicines Agency issued reporting guidance for PBPK analyses in regulatory submissions and is extending that framework to mechanistic models including PBPK and quantitative systems pharmacology (EMA, 2018; FDA, 2016). Verification, validation and uncertainty quantification therefore determine what a model can support; they do not convert every simulation into regulatory evidence (Viceconti et al., 2021).

xISCT was designed around traceability. It combines harmonized molecular data, probabilistic protein states, a Human Protein Network, an ensemble of inferred network solutions, supervised phenotype calibration, simplified PBPK exposure and predefined statistical controls (Figure 1). The platform estimates patient-specific target perturbation, propagates it through alternative xModels and quantifies downstream effects at disease effectors, pathways and responder-associated proteins. Meta-algorithms make the inference and calibration logic explicit, while ensemble recurrence prevents a single optimized graph from carrying the conclusion.

We applied xISCT to escitalopram 10 and 20 mg/day, citalopram 40 mg/day and duloxetine 60 mg/day. The paired escitalopram doses test internal stability; citalopram tests whether a closely related SSRI preserves a common topology while redistributing signal; and duloxetine tests a serotonergic–noradrenergic entry. The primary objective was to compare complete and individual mechanisms of action. Secondary objectives were to quantify target and effector sensitivity, characterize baseline and treatment-emergent differences between higher and lower responders and assess whether model-prioritized proteins agree with experimental or clinical literature. Aggregate response outputs were retained only as model-behaviour controls.

3 Methods




Figure 1. Conceptual workflow of the xISCT study. Public molecular profiles and a prespecified study template were mapped to a Human Protein Network and used to construct separate Diseased, Simulated Diseased and Healthy cohorts. Patient-specific exposure perturbed 100 semisupervised xModels. Outputs included phenotype consistency, compared and individual mechanisms of action and responder heterogeneity. VHT, virtual human twin; PBPK, physiologically based pharmacokinetics.

3.1 Study design and interventions

The study used four longitudinal virtual-treatment arms with baseline and final states (Table 1). Escitalopram 10 mg/day (ESCIT10), escitalopram 20 mg/day (ESCIT20) and citalopram 40 mg/day (CITAL40) were simulated for 8 weeks (56 once-daily doses). Duloxetine 60 mg/day (DULOX60) was simulated for 9 weeks (63 once-daily doses). The same eligible patient could be evaluated under more than one treatment. The enrolment count therefore denotes unique virtual patients, whereas arm denominators denote patient-treatment observations and must not be summed as independent enrolment.

Table 1. Virtual interventions and analysis frame

Arm	Regimen	Duration	Final analysed N	Role in comparison
ESCIT10	Escitalopram 10 mg once daily	8 weeks	508	Dose-pair stability
ESCIT20	Escitalopram 20 mg once daily	8 weeks	500	Dose-pair stability
CITAL40	Citalopram 40 mg once daily	8 weeks	509	Related SSRI
DULOX60	Duloxetine 60 mg once daily	9 weeks	507	SNRI comparator

Note. N is the patient–treatment denominator reported for the primary response analysis.

3.2 Virtual population and molecular cohorts

POPU_DEPRESSION represented adult outpatients with MDD in the United States: 33.8% men, 66.2% women and an approximate age range of 18.5–64.5 years. Diseased contained repository patients annotated with depression. Simulated Diseased was created from initially healthy patients by setting selected disease effectors in the disease direction and applying one xModel propagation pulse. Healthy patients remained an independent reference for relative protein activation or inhibition. Five hundred patients were requested per cohort; the repository yielded 154 Diseased, 500 Simulated Diseased and 500 Healthy patients, or 1,154 unique individuals.

3.3 Probabilistic molecular states

Public molecular measurements were cross-normalized, discretized into ten bins and represented as protein-specific probability distributions. Values already on a compatible log scale were standardized directly; other values were log transformed before standardization.

$$z = (x - \mu)/\sigma \quad \text{or} \quad z = (\log_2(x + \varepsilon) - \mu_{\log})/\sigma_{\log} \quad (1)$$

Cross-normalization; ε prevents numerical instability before logarithmic transformation.

Relative activation or inhibition was calculated against the independent Healthy reference on a signed log-odds scale. Probabilities were clipped to $[10^{-6}, 1-10^{-6}]$.

$$\text{logit}(p) = \ln[p/(1-p)] \quad (2)$$

$$\Delta L = \text{logit}(p_{\text{evaluated}}) - \text{logit}(p_{\text{healthy}}); \quad s = 2 \cdot \text{logistic}(\Delta L) - 1 \quad (3)$$

The signed state s lies between -1 (relative inhibition) and $+1$ (relative activation).

3.4 Human Protein Network and xModel ensemble

The Human Protein Network comprised approximately 14,700 proteins and 270,000 interactions. Approximately 40,300 links (15%) had known functional directions and were fixed. Unknown link signs were inferred under 1.05 million individual constraints consolidated into 461 grouped rules. For each candidate signed network, constraint-specific inputs were propagated through three iterative hyperbolic-tangent cycles. Candidate configurations were evaluated using an aggregate signed constraint loss that jointly accounts for agreement with the expected effector directions and the magnitude of the propagated response. New configurations were generated by flipping subsets of unknown edge signs and were accepted using a temperature-controlled Metropolis-type rule, while the best-so-far configuration was retained. All 100 available xModels were retained, and conclusions were based on recurrence across the ensemble rather than on a single optimized graph.

$$L(s) = -N^{-1} \sum_{i=1}^N \ln p_i(s), \quad p_i(s) = \text{clip}((\text{clip}[d_i(s)a_i(s), -1, 1] + 1)/2, \varepsilon, 1) \quad (4)$$

Aggregate signed constraint loss used to evaluate candidate network-sign configurations.

For constraint i , $d_i(s)$ measures the directional agreement between the final propagated signals and the expected signs of its predefined effectors. It is calculated as the sum of the expected-sign-weighted effector signals divided by the larger of their total absolute signal and epsilon. The term $a_i(s)$ is the maximum absolute response among those effectors, thereby penalizing configurations that produce the correct direction but negligible signal. The implementation uses $\epsilon = 10^{-6}$. The resulting $p_i(s)$ is a bounded, probability-like constraint-compliance score and is not a calibrated probability of a biological or clinical outcome. The complete constrained-inference sequence is summarized in Table 2.

Table 2. Meta-algorithm for constrained xModel inference

Step	Operation	Implementation
1	Initialize	Known edge signs fixed; unknown edge signs randomized
2	Select	Use the active subset of training constraints
3	Propose	Flip the signs of a subset of unknown edges
4	Evaluate	Three propagation cycles; aggregate signed loss (Eq. 4)
5	Accept/reject	Temperature-controlled Metropolis-type rule based on delta loss
6	Update	Retain the current state and the best-so-far configuration
7	Bias	Penalize repeatedly rejected flips and adapt proposal size
8	Explore	Parallel population strategy
9	Aggregate	Canonical evaluation on all constraints; retain the xModel ensemble

3.5 Supervised phenotype calibration

Disease and clinical-variable effectors were filtered by tissue availability, validation and expected direction, assembled into candidate clusters and evaluated across supervised scenarios. Twenty-four metrics were tested. The fixed solution used `posNegBalance_x_media`, retained three calibrated entities (DEPRESSION, and variables V1 and V2, both other indirect measures about depression), fifteen selected proteins. The chosen configuration was frozen before final MoA, pathway, cohort-concordance and responder analyses. Downstream outputs therefore interpret a calibrated model and are not an external validation of that model.

Table 3 records the prespecified calibration and freezing sequence used before the final simulations.

Table 3. Meta-algorithm for supervised phenotype calibration

Step	Operation	Purpose
1	Define targets	Pathology and clinical variables
2	Filter effectors	Tissue, validation and measurability
3	Assign direction	Expected phenotype relationship
4	Generate clusters	Alternative protein combinations
5	Evaluate scenarios	Test simulated treatments
6	Check ordering	Compare with prespecified ranking
7	Select solution	Metric, clusters and proteins

Step	Operation	Purpose
8	Freeze and simulate	No post-hoc target redefinition
9	Validate outputs	Coverage, robustness, concordance and limits

3.6 Patient-specific exposure

Exposure was simulated with a simplified one-compartment repeated-dose PBPK model using a 0.01-hour integration step. For each patient, scheduled doses contributed to the concentration-time profile; effective concentration incorporated unbound fraction and target-tissue availability. The resulting bounded exposure factor scaled every drug–target perturbation.

$$C(t) = \Sigma k Dose_k / Vd \cdot e^{(-ke(t - tk))} \cdot 1(t \geq tk) \tag{5}$$

$$E_{patient} = (1/T) \int_0^T C_{eff}(t) / [C_{eff}(t) + Potency] dt \tag{6}$$

$$Signal_{patient, target} = E_{patient} \cdot IntrinsicDrugEffect_{target} \tag{7}$$

The patient exposure factor is bounded between 0 and 1; the intrinsic effect preserves target sign.

Figure 2 presents the factor that directly enters network propagation. Raw concentration profiles are retained as Supplementary Figure S1 because PBPK is an upstream scaling component rather than the principal biological result.




Figure 2. Patient-level exposure-factor distributions used to scale drug–target perturbation. Solid lines represent Simulated Diseased and dashed lines Diseased patients. The bounded exposure factor, rather than the raw PBPK concentration curve, enters network propagation.

3.7 Mechanism-of-action, pathway and responder analyses

The complete mechanism propagated exposure-weighted perturbations for three cycles through all available proteins. Simplified graphs connected targets to selected depression effectors through paths

of up to three links. Proteins were classified as sources, sinks, transmitters or low-signal/ambiguous nodes. Target importance was assessed by removing one target at a time while holding exposure, xModels and propagation cycles fixed. Sensitivities are perturbational dependencies and were not interpreted as independent causal fractions.

Gene set enrichment analysis ranked all proteins with propagated signal and paired measurement; no significance or top-N filter was applied before enrichment (Subramanian et al., 2005). FDR correction controlled multiple testing (Benjamini and Hochberg, 1995). For responder analysis, only simulated responders were retained and the upper and lower 25% were compared. This threshold maximized a separation score of 20.28 while preserving at least 37 patients per tail. Post-treatment contrasts used within-patient final-minus-baseline change; baseline analyses used pretreatment molecular states. All complete exported tables were analysed, rather than only report excerpts.

$$\Delta Contrast = median(\Delta protein | higher responders) - median(\Delta protein | lower responders) \quad (8)$$

Negative values indicate greater inhibition or less activation in higher responders.

3.8 Statistical controls and interpretation boundaries

Rank-sum and Fisher exact tests were adjusted by the Benjamini–Hochberg procedure. Molecular results at $q < 0.10$ were exploratory; $q < 0.05$ provided stronger adjusted evidence. Controls evaluated patient retention, phenotype objectives, clinical-variable coverage and order, cluster robustness and agreement between Diseased and Simulated Diseased sources. Aggregate response estimates were retained to test plausible model behaviour, not to rank treatments. The study did not establish safety, clinical superiority or prescribing recommendations.

4 Results

4.1 Population flow and quality-control boundary

The Simulated Diseased cohort retained 404/498 (81.1%), 400/498 (80.3%), 415/498 (83.3%) and 408/498 (81.9%) patients under ESCIT10, ESCIT20, CITAL40 and DULOX60, respectively. The Diseased cohort retained 104/154 (67.5%), 100/154 (64.9%), 94/154 (61.0%) and 99/154 (64.3%). The minimum 61.0% exceeded the prespecified 48.6% retention threshold. Cluster robustness and cross-cohort concordance passed: no low-robustness row and no opposite direction between disease sources were reported.

The central limitation was phenotype calibration. V1 and V2 achieved 100% coverage, prediction correlation $r=0.97$ and RMSE 5.29 percentage points, but the expected clinical ordering was not reproduced (order correlation $r=0.47$; validity status REVISAR). Thirty-three of 102 phenotype-objective controls also required review (Table 4). Model responder proportions were 30.12%, 34.00%, 29.47% and 32.15%, and calibrated response values were 62.01%, 64.13%, 61.60% and 62.56%, respectively. These overlapping values are descriptive controls and cannot support comparative efficacy.

Table 4. Population retention and model-control results

Domain	Result
Retention	61.0–83.3% across arms/cohorts
xModel ensemble	100 models; mean constraint loss 0.14
Clinical variables	coverage 100%; $r_{\text{prediction}}=0.97$; RMSE=5.29 pp
Phenotype objectives	33/102 unmet
Cohort agreement	No opposite directions

4.2 Target sensitivity distinguishes pharmacological entry

Escitalopram was SERT-dominant at both doses: SERT accounted for 72.74% and 71.85% of relative network sensitivity, sigma-1 for 12.83% and 13.30%, and DRD2, HTR1B and HTR1A for approximately 3.2–3.9% each. Citalopram retained SERT as the largest dependency (43.90%) but redistributed sensitivity towards sigma-1 (28.56%), DRD2 (5.79%), HTR1B (5.25%), H1 (5.05%) and HTR1A (4.58%). Duloxetine was dominated by SERT (91.46%), with smaller NET (7.17%) and DAT (1.37%) contributions (Figure 3). These values quantify dependence within the fitted network and should not be equated with receptor occupancy or clinical contribution.




Figure 3. Leave-one-target-out global network sensitivity. The escitalopram doses show a stable SERT/sigma-1 distribution; citalopram redistributes sensitivity across a broader receptor set; duloxetine is SERT-dominant with smaller NET and DAT components.

4.3 Compared mechanisms define a conserved SSRI topology and a compact duloxetine network

ESCIT10, ESCIT20 and CITAL40 each generated 254 informative MoA proteins: 188 functional transmitters, 63 low-signal/ambiguous proteins and three functional sinks. Their protein sets were identical (pairwise Jaccard=1.00), despite the treatment-specific weights shown by target sensitivity. Duloxetine generated 43 proteins: 30 transmitters, 11 low-signal/ambiguous proteins and two sinks. Thirty-eight proteins were shared across all four arms; duloxetine added five proteins absent from the SSRI topology (Figure 4). Pairwise Jaccard similarity between duloxetine and each SSRI arm was 0.147.

VAMP2 and SCAMP2 were functional sinks in every arm; GHSR was the additional sink in the three SSRI arms. Conserved sinks identify points at which a network with different pharmacological inputs repeatedly attenuated signal. The complete comparative network integrates these relationships in one view (Figure 5). Serum amyloid P has been shown experimentally to promote VAMP2-mediated SERT membrane insertion and was associated with SSRI response in MDD, giving independent plausibility to the VAMP2/SERT interface (Su et al., 2023). The sink classification itself remains a model prediction.




Figure 4. Size and four-arm overlap of informative mechanism-of-action protein sets. The three SSRI arms share the same 254-protein topology; duloxetine retains 38 of its 43 proteins in the four-arm core.




Figure 5. Complete comparative mechanism-of-action network for escitalopram 10 mg/day, escitalopram 20 mg/day, citalopram 40 mg/day and duloxetine 60 mg/day. Treatment colours identify entry-specific routes, whereas grey elements are shared. The figure is displayed on one A4 landscape page to preserve the visual topology.

4.4 Individual mechanisms preserve topology while changing signal weights

Individual graphs clarify how the same SSRI protein set can carry different quantitative contributions. The escitalopram maps are topologically identical and differ through dose- and patient-specific exposure, supporting an internal stability check (Figure 6a). Citalopram uses the same 254-node set but gives greater weight to sigma-1, H1 and additional receptors, demonstrating redistribution within a conserved topology (Figure 6b). Duloxetine enters through SERT, NET and DAT but propagates through a much smaller informative network (Figure 6c). Network breadth is therefore not a measure of clinical efficacy; it is a property of the fitted disease-conditioned propagation.




Figure 6a. Individual mechanism of action for escitalopram 10 mg/day. The escitalopram 20 mg/day graph has the same topology; quantitative amplitudes remain exposure- and patient-specific.




Figure 6b. Individual mechanism of action for citalopram 40 mg/day. The protein set is identical to the escitalopram topology, but target and downstream weights are redistributed.




Figure 6c. Individual mechanism of action for duloxetine 60 mg/day. The compact SERT/NET/DAT entry converges on VMAT2 and FGF1 and retains VAMP2/SCAMP2 as functional sinks.

4.5 VMAT2–FGF convergence and citalopram-specific CRH/FGFR1 amplification

VMAT2 and FGF1 were the dominant selected effectors in all arms. VMAT2 activation ranged from 47.10 (DULOX60) to 55.20 (ESCIT20), and FGF1 from 46.97 to 55.00. FGFR1 was most strongly activated by CITAL40 (6.58 versus 1.87, 2.02 and 0.21), while CRH inhibition was also strongest under CITAL40 (–17.76 versus –4.74, –5.12 and approximately zero). CRHR2 signals were small (0.00–0.88) (Figure 7 and Table 5).

VMAT2 is biologically credible because reduced Vmat2 expression produces depressive-like phenotypes in mice that respond to antidepressants (Fukui et al., 2007). The FGF system is dysregulated in human MDD cortex, with SSRI exposure attenuating some transcript changes (Evans et al., 2004), and recent experimental work linked a dentate-gyrus FGFR1–Notch–BDNF axis to neurogenesis and antidepressant effects (Shin et al., 2025). These studies support monoamine-storage and trophic-signalling domains; they do not validate the exact numerical amplitudes or citalopram specificity predicted here.

Signed propagated signal at selected depression effectors				
	ESCIT10	ESCIT20	CITAL40	DULOX60
VMAT2	53.69	55.20	51.31	47.10
FGF1	53.51	55.00	50.75	46.97
FGFR1	1.87	2.02	6.58	0.21
CRH	-4.74	-5.12	-17.76	-0.00
CRHR2	0.15	0.16	0.88	0.00

Positive values indicate activation; negative values indicate inhibition.

Figure 7. Signed propagated signal at five selected depression effectors. VMAT2 and FGF1 are shared high-amplitude outputs; citalopram produces the largest FGFR1 activation and CRH inhibition.

Table 5. Effector-level mechanistic comparison

Effector	ESCIT10	ESCIT20	CITAL40	DULOX60	Interpretation
VMAT2	53.69	55.20	51.31	47.10	Shared monoamine-storage output
FGF1	53.51	55.00	50.75	46.97	Shared trophic output
FGFR1	1.87	2.02	6.58	0.21	Citalopram-amplified
CRH	-4.74	-5.12	-17.76	-0.00	Citalopram-amplified inhibition
CRHR2	0.15	0.16	0.88	0.00	Low-amplitude output

Note. Values are signed effect per patient from the stored report; all reported p-values were <0.01.

4.6 Complete responder tables reveal a broad concordant response core

The complete post-treatment tables contained 493, 2,014, 1,319 and 1,235 proteins at exploratory $q < 0.10$ for ESCIT10, ESCIT20, CITAL40 and DULOX60, respectively; 165, 1,822, 977 and 1,059 also passed $q < 0.05$. Four hundred thirty-four proteins were present in all arms with the same higher-versus-lower-responder direction (Figure 8). This shared core included BLNK, serum amyloid P, claudin-5, WNT1, KCNJ8, CD44 and HSPG2. Every highlighted contrast was negative, indicating greater inhibition or less activation in higher responders.

The strongest treatment-conditioned entries were BLNK/HSPG2/PTGER2/SNAP23/GRIK3 in ESCIT10; WNT1/claudin-5/serum amyloid P/BLNK/KCNJ8/CD44 in ESCIT20; CD47/KCNJ8/SLC2A5/CD44/NPPA in CITAL40; and striatin/WNT1/KCNJ8/CCR9/CD44/HSPG2/GRM4 in DULOX60. These are statistical contrasts among simulated responders, not validated biomarkers or treatment-selection rules.

Several signals have external support at the domain level. Higher serum amyloid P has been associated with better escitalopram response and experimentally alters SERT localization (Su et al., 2023). Social stress reduces claudin-5, increased blood–brain barrier permeability and depression-like behaviour; chronic antidepressant treatment rescued claudin-5 loss in mice, while human MDD tissue showed reduced expression (Menard et al., 2017). Duloxetine-response microRNAs regulate MAPK/WNT genes across a clinical trial, replication cohorts and experimental systems (Lopez et al., 2017). CD47 protein is elevated in synaptic fractions from MDD cortical tissue, connecting the CITAL40 signal to neuronal regulation of microglia, although not directly to treatment response (Scheepstra et al., 2023). KCNJ8, HSPG2, BLNK and striatin remain prioritized hypotheses without equally direct antidepressant-response evidence.




Figure 8. Full-table responder analysis. Left, proteins meeting exploratory $q < 0.10$ in each arm and the 434-protein direction-concordant core. Right, higher-versus-lower-responder contrasts for four recurrent candidates; negative values indicate greater inhibition or less activation in higher responders.

Table 6 separates literature-supported convergence from model-generated hypotheses and specifies a direct validation route for each prioritized candidate.

Table 6. Evidence level and direct validation route for prioritized responder proteins

Candidate	Observed pattern	External evidence	Status / next test
SAP (P02743)	Negative contrast in all arms; strongest ESCIT20	Clinical/experimental link to SERT trafficking and SSRI response	Targeted plasma assay plus membrane-SERT measurement
CLDN5 (O00501)	Negative contrast in all arms; strongest ESCIT20	Stress, BBB and antidepressant-rescue evidence	Longitudinal endothelial/BBB markers
WNT1 (P04628)	Negative contrast in all arms	Duloxetine-response miRNAs regulate MAPK/WNT genes	Paired transcript/protein validation
CD47 (Q08722)	Top CITAL40 contrast	Elevated synaptic CD47 in MDD cortex	Microglial phagocytosis and response study
KCNJ8 / HSPG2	Recurrent multi-arm contrasts	Indirect or limited response evidence	Perturbation and prospective replication

Note. Evidence tiers calibrate language; biological concordance does not constitute prospective validation.

4.7 Higher- and lower-responder MoAs localize the statistical contrasts

The same general treatment MoA was overlaid with response-associated proteins for the upper and lower response tails. ESCIT10 compared 38 versus 38 responders, ESCIT20 42 versus 42, CITAL40 37 versus 37 and DULOX60 40 versus 40. Red nodes indicate negative higher-minus-lower responder contrasts, green nodes indicate positive contrasts, and grey nodes represent shared or non-differential mechanisms. Dashed green nodes indicate statistically supported positive-response candidates outside the retained paths. The overlays connect the response tables to treatment-specific topology (Figures 10a–10d).

ESCIT10 localized HSPG2, PTGER2, BLNK, SNAP23 and GRIK3 to the general SERT-centred mechanism. ESCIT20 localized WNT1, claudin-5, serum amyloid P, BLNK and KCNJ8, creating the strongest literature-supported SSRI responder overlay through SAP/SERT trafficking and claudin-5 neurovascular biology. CITAL40 emphasized CD47, KCNJ8, SLC2A5 and CD44 within a receptor-distributed SSRI network.

DULOX60 emphasized WNT1, striatin, KCNJ8, CCR9, CD44, HSPG2 and GRM4 in its compact SERT/NET/DAT topology. The figures support mechanistic localization; the overlaid proteins cannot be assumed necessary or sufficient for response.

4.8 Baseline profiles are treatment-conditioned and absent for citalopram

Pretreatment contrasts differed markedly in scale. ESCIT10 contained 35 proteins at $q < 0.10$ (one at $q < 0.05$); the leading candidates were mitochondrial COX7A2 ($q = 0.040$), cytochrome c ($q = 0.072$) and the apoptosis regulator HRK ($q = 0.080$). ESCIT20 contained 598 proteins at $q < 0.10$ and 344 at $q < 0.05$, including CCN1, BIRC5, BEST1, FZD9, SNAP23, FAAH and DLL1. No CITAL40 baseline protein passed $q < 0.10$, which prevents a treatment-specific pretreatment molecular profile for that arm. DULOX60 contained 588 proteins at $q < 0.10$ and 320 at $q < 0.05$; SREBF1 was lower in higher responders ($q = 0.0047$), PCSK9 higher ($q = 0.0048$) and ELK1 lower ($q = 0.0074$). Because the same underlying population can appear in several treatment arms, differences reflect treatment-conditioned eligibility, exposure and response stratification rather than four independent cohorts.

Demographic signals were exploratory. ESCIT20 higher responders were older by a median 10 years ($q = 0.09$). A combined older-age-plus-male feature reached $q = 0.01$ at population level and under duloxetine, while individual age and sex comparisons did not survive FDR in the other arms. These variables contribute to simulated PBPK exposure and cannot be interpreted as independent clinical predictors.

Table 7 provide the arm-level synthesis of responder-tail sizes, post-treatment candidates, baseline priorities and pathway findings. Figure 9 integrates these findings by linking each treatment to selected molecular signatures associated with lower response and demographic characteristics associated with higher response. Solid links represent exploratory associations meeting FDR $q < 0.10$, whereas dashed links indicate descriptive demographic associations.

Table 7. Treatment-conditioned responder and baseline profiles

Arm	Higher/lower tail	Post-treatment priorities	Baseline priorities	Pathway result
ESCIT10	38 / 38	BLNK, HSPG2, SNAP23, GRIK3, PTGER2	COX7A2; cytochrome c; HRK	Innate/adaptive immune Reactome term, $q = 0.07$
ESCIT20	42 / 42	WNT1, CLDN5, SAP, KCNJ8, CD44	CCN1, BIRC5, BEST1, FZD9, SNAP23	None at $q < 0.10$
CITAL40	37 / 37	CD47, KCNJ8, SLC2A5, CD44, NPPA	None at $q < 0.10$	NOTCH-labelled Reactome terms, $q \approx 0.10$
DULOX60	40 / 40	STRN, WNT1, KCNJ8, CCR9, CD44, GRM4	SREBF1, PCSK9, ELK1	None at $q < 0.10$

Note. Pathway overrepresentation is exploratory and based on differential responder proteins.




Figure 9: Treatment-associated profiles of higher and lower responders. The alluvial diagram links each treatment to selected demographic and molecular features associated with response extremes. Solid links indicate exploratory FDR $q < 0.10$, whereas dashed links show descriptive demographic context; link width reflects visual prioritization, not patient counts.

4.9 Pathway enrichment is sparse and exploratory

Only ESCIT10 and CITAL40 produced pathway rows below the exploratory $q < 0.10$ threshold. ESCIT10 enriched one broad Reactome innate/adaptive immune-response term (25 hits, $q = 0.07$). CITAL40 enriched six closely related NOTCH-labelled Reactome terms at $q \approx 0.10$, including NOTCH1 mutant/cancer annotations and NOTCH2 signalling. The database labels should not be interpreted as evidence of cancer biology in the simulated patients; the shared hit set indicates NOTCH-related signalling among response-associated proteins. No pathway passed $q < 0.10$ for ESCIT20 or DULOX60. These sparse results do not negate protein-level signals, but they prevent a strong pathway-level treatment narrative.




Figure 10. Compact mechanisms of action distinguishing higher and lower simulated responders. Panels show (a) escitalopram 10 mg/day (38 versus 38 patients), (b) escitalopram 20 mg/day (42 versus 42), (c) citalopram 40 mg/day (37 versus 37) and (d) duloxetine 60 mg/day (40 versus 40). Drug targets, selected effectors, response-associated proteins and their shortest connecting paths were retained; terminal proteins with one incoming and no outgoing link were omitted for clarity. Red nodes indicate negative higher-minus-lower response contrasts, green nodes positive contrasts and grey nodes shared or non-differential mechanisms. Dashed green nodes represent statistically supported positive-response candidates outside the retained paths. Arrows indicate activation and blunt ends inhibition. The SSRI networks highlight synaptic, neurovascular, inflammatory and metabolic components—including HSPG2, PTGER2, SNAP23, GRIK3, WNT1, CLDN5, SAP, CD44, SLC2A5, NPPA and KCNJ8—whereas duloxetine presents a more restricted SERT–VAMP2-centred topology involving HSPG2, STRN, WNT1, KCNJ8 and GRM4. KRT10, CHST15, RASSF1 and PTPN1 were the strongest positive contrasts in their respective arms. These patterns are mechanistic hypotheses rather than validated predictive biomarkers.

5 Discussion

5.1 A stable SSRI topology can carry different pharmacological weights

The three SSRI arms produced the same 254-protein topology, yet citalopram redistributed leave-one-target-out sensitivity from SERT towards sigma-1, H1, dopamine and serotonin receptors. This separation is methodologically important: topology records which informative routes are available, whereas perturbation weights record how a particular intervention uses those routes. The paired escitalopram

doses therefore act as an internal stability test, and the citalopram arm shows that a related drug can preserve topology while changing quantitative propagation.

5.2 Duloxetine separates target catalogue from network breadth

Duloxetine entered through SERT, NET and DAT but produced only 43 informative MoA proteins, of which 38 belonged to the four-arm core. Its target catalogue is therefore broader than its retained disease-conditioned network. This illustrates why direct targets alone are insufficient for systems-level interpretation: NET contributed 7.17% to global sensitivity even though the selected depression effectors were dominated by the shared VMAT2–FGF output. A compact network is not weaker clinical efficacy; it is a falsifiable prediction about constrained propagation in the calibrated molecular state.

5.3 VMAT2, FGF and CRH define a convergent effector layer

The clearest convergence occurred downstream. VMAT2 and FGF1 were strongly activated in every arm, while citalopram amplified FGFR1 and CRH effects. Reduced Vmat2 expression produces depressive-like phenotypes responsive to antidepressants in mice (Fukui et al., 2007). Human post-mortem work identified FGF-system dysregulation in MDD (Evans et al., 2004), and an experimental FGFR1–Notch–BDNF axis was recently linked to neurogenesis and antidepressant effects (Shin et al., 2025). CRH connects the model to hypothalamic–pituitary–adrenal stress biology, but the treatment-specific signed pattern lacks direct validation and should be tested rather than narrated as established mechanism.

5.4 Patient profiles add explanation beyond endpoint comparisons

A conventional trial can establish whether symptom scores differ on average; the present framework decomposes the same response label into a baseline molecular position and a treatment-emergent trajectory. The 434-protein concordant response core argues against entirely separate responder biology for each drug. Arm-conditioned deviations then identify where a specific perturbation interacts with that shared substrate. This distinction is the central added value of the virtual-twin analysis: it proposes molecular reasons for simulated heterogeneity while preserving the uncertainty around every candidate.

The best-supported convergence is the ESCIT20 overlay of serum amyloid P and claudin-5. SAP has direct clinical and experimental links to SSRI response and SERT trafficking (Su et al., 2023), while claudin-5 has causal preclinical and human-tissue support for stress susceptibility and antidepressant rescue (Menard et al., 2017). WNT biology is supported by a duloxetine trial and functional experiments (Lopez et al., 2017). CD47 is linked to altered neuron–microglia regulation in MDD cortex (Scheepstra et al., 2023). These correspondences validate biological domains, not the model's exact direction or predictive performance. BLNK, KCNJ8, HSPG2, striatin and several baseline proteins remain valuable precisely because the model makes them explicit candidates for prospective validation.

5.5 Methodological contribution

xISCT integrates known and inferred biology with explicit control logic. Known network directions remain fixed while unknown functions are inferred through a nine-step xModel meta-algorithm; 100 solutions reduce dependence on one graph and preserve heterogeneity to capture the most representative possible explanations. A separate calibration meta-algorithm selects and freezes fifteen proteins before final mechanistic analysis. PBPK translates patient covariates into a bounded perturbation, and full-table statistics distinguish recurrent signals from report excerpts. Equations 1–8 expose the normalization, state comparison, optimization, exposure and responder contrasts.

5.6 Study Scope and Interpretative Considerations

First, this is a simulated study. Preservation of molecular constraints and cross-cohort agreement does not replace prospective clinical validation. Second, the expected treatment order was not reproduced, both clinical variables were flagged as REVISAR, and 33 of 102 phenotype-objective checks were unmet. Consequently, the simulated response percentages cannot be interpreted as evidence of comparative

efficacy. However, the confidence intervals around the treatment-level estimates overlapped substantially, indicating that the apparent numerical differences between arms were small relative to their uncertainty. Third, the molecular repository was heterogeneous and availability-limited; only 154 directly Diseased patients were available. The inclusion of independent Diseased and Simulated Diseased cohorts for each treatment arm provided an internal robustness check, with no opposite directions observed between cohort sources, but it does not fully overcome the limited size of the directly Diseased cohort. Fourth, approximately 85% of network-link directions were inferred rather than known. An ensemble of 100 models and fixed biological restrictions reduces, but does not eliminate, structural uncertainty. Fifth, the PBPK component was simplified and was used to derive relative patient-specific exposure scaling rather than to provide compound-specific clinical pharmacokinetic validation. This level of modelling was appropriate for its defined role in the present study, but it does not support detailed conclusions about the drugs' clinical pharmacokinetic behaviour. Sixth, responder analyses used an exploratory $q < 0.10$ threshold and compared the upper and lower response tails among simulated responders. Seventh, pathway enrichment was sparse, and several database annotations were not specific to depression. Finally, the study does not establish clinical safety, clinical superiority or prescribing recommendations.

5.7 Validation roadmap

Validation should proceed in tiers. First, targeted proteomic or transcriptomic assays in stored longitudinal samples can test direction and treatment association for SAP, claudin-5, WNT1, CD47, KCNJ8, HSPG2 and striatin. Second, neuronal, endothelial, astrocytic and immune-cell perturbation can test causal positions for VMAT2, FGF1/FGFR1, CRH and the conserved VAMP2/SCAMP2 sinks. Third, prospective studies should pre-specify exposure measurement, baseline candidates and higher-versus-lower response contrasts. The strict test is not whether a candidate is associated with MDD, but whether it adds reproducible information about response under a defined treatment and measured exposure.

6 Conclusions

xISCT enabled the molecular simulation and head-to-head mechanistic comparison of three antidepressants administered through four treatment regimens. Starting from public expression-derived molecular profiles, drug–target annotations, biological constraints and a Human Protein Network, the platform reconstructed a conserved 254-protein topology for the SSRI interventions, a more compact 43-protein mechanism for duloxetine and a shared downstream layer dominated by VMAT2 and FGF signalling. The two escitalopram doses reproduced the same network topology, citalopram redistributed the relative contribution of several targets within that shared topology, and duloxetine produced a distinct SERT/NET/DAT entry architecture. These results show that the platform can recover recognizable properties of antidepressant pharmacology while also describing how those properties propagate through a disease-conditioned molecular system.

A principal contribution of this work is the operational connection between molecular states and clinically defined phenotypes. Clinical definitions were translated into supervised molecular objectives, allowing patient-specific drug perturbations to be related to simulated phenotype changes, response strata and treatment-conditioned molecular profiles. This connection provides a means of moving beyond conventional comparisons of aggregate clinical endpoints towards an explanation of the proteins and network routes that may contribute to heterogeneous treatment responses. Nevertheless, the failure to reproduce the complete prespecified clinical treatment order and the unresolved phenotype-objective checks limit the interpretation of the simulated response percentages. The present results therefore support mechanistic interpretation and hypothesis generation, rather than claims of comparative clinical efficacy.

Analysis of the complete responder tables identified a 434-protein direction-concordant core shared across all four interventions, together with treatment-conditioned signatures that distinguished higher

and lower simulated responders. Several prioritized proteins and pathways were independently supported by published experimental or clinical observations, including serum amyloid P-dependent regulation of SERT trafficking, claudin-5-associated neurovascular biology, VMAT2-dependent monoamine storage, FGF/FGFR1 signalling and WNT-related mechanisms. This external concordance strengthens the biological plausibility of the reconstructed networks without validating their precise quantitative effects. Other recurrent candidates—including BLNK, KCNJ8, HSPG2, striatin and selected baseline proteins—represent potentially novel hypotheses that warrant direct molecular investigation and prospective clinical replication.

More broadly, the study illustrates how biocomputational approaches such as xISCT can integrate public molecular data, systems-level biological knowledge, pharmacological exposure and clinically meaningful definitions within a traceable virtual-trial framework. Such approaches are not replacements for experimental studies or randomized clinical trials. Rather, they can help explain disease and drug mechanisms, prioritize candidate biomarkers, identify sources of responder heterogeneity and guide the design of more focused experimental and clinical studies. By narrowing the space of plausible mechanisms and making patient-level molecular hypotheses explicit before prospective testing, virtual-human-twin and in silico clinical-trial platforms may contribute to more informative clinical development strategies and to the identification or optimization of future therapeutic interventions.

7 Data, code and reporting statements

Data provenance. All project numbers, model outputs, tables and figures in this manuscript were derived from the DEPRESSION10 xISCT report, its complete exported tables and the original project image/DOT assets. External publications and regulatory documents were used only for epidemiological context, methods positioning and biological interpretation; no external dataset was added to the project analysis.

Code availability. xISCT is independently developed proprietary computational software. Its source code and trained model objects are not publicly distributed. The methodological procedures are described in this manuscript to facilitate independent implementation and evaluation by other researchers.

Ethics. This work is an in silico study based on de-identified public molecular resources assembled in the project repository and does not report a new intervention in human participants. The final submission should include the source-data and ethics statements required by the contributing datasets.

Funding. This research received no external funding.

Competing interests: The authors declare no competing interests.

8 Author contributions

The computational software underlying this study was independently conceived, designed, developed and implemented by J.M.M., drawing on expertise in computational systems biology, network pharmacology and mechanistic modelling, and using personal computational resources. E.R. contributed operational guidance, systematic monitoring and validation of computational workflows, supporting consistency and methodological robustness. X.D. provided scientific supervision and expert guidance in biological and molecular interpretation, critical discussion and manuscript review. The source code is not distributed as part of this publication.

9 References

1. Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Stat Soc Series B*. 1995;57:289–300.

2. Bymaster FP, Dreshfield-Ahmad LJ, Threlkeld PG, et al. Comparative affinity of duloxetine and venlafaxine for serotonin and norepinephrine transporters in vitro and in vivo. *Neuropsychopharmacology*. 2001;25:871–880. PMID:11750180.
3. Chan MK, Cooper JD, Bot M, et al. Blood-based immune-endocrine biomarkers of treatment response in depression. *J Psychiatr Res*. 2016;83:249–259. doi:10.1016/j.jpsychires.2016.08.020.
4. Cipriani A, Furukawa TA, Salanti G, et al. Comparative efficacy and acceptability of 21 antidepressant drugs for acute treatment of adults with major depressive disorder. *Lancet*. 2018;391:1357–1366. doi:10.1016/S0140-6736(17)32802-7.
5. Detke MJ, Lu Y, Goldstein DJ, Hayes JR, Demitrack MA. Duloxetine, 60 mg once daily, for major depressive disorder: a randomized double-blind placebo-controlled trial. *J Clin Psychiatry*. 2002;63:308–315. doi:10.4088/JCP.v63n0407.
6. European Medicines Agency. Guideline on the reporting of physiologically based pharmacokinetic modelling and simulation. EMA/CHMP/458101/2016. 2018; effective 1 July 2019.
7. Evans SJ, Choudary PV, Neal CR, et al. Dysregulation of the fibroblast growth factor system in major depression. *Proc Natl Acad Sci USA*. 2004;101:15506–15511. doi:10.1073/pnas.0406788101.
8. Fukui M, Rodriguiz RM, Zhou J, et al. Vmat2 heterozygous mutant mice display a depressive-like phenotype. *J Neurosci*. 2007;27:10520–10529. doi:10.1523/JNEUROSCI.4388-06.2007.
9. Gadad BS, Jha MK, Grannemann BD, Mayes TL, Trivedi MH. Proteomics profiling reveals inflammatory biomarkers of antidepressant treatment response: findings from the CO-MED trial. *J Psychiatr Res*. 2017;94:1–6. doi:10.1016/j.jpsychires.2017.05.012.
10. Gutiérrez-Casares JR, Quintero J, Jorba G, Junet V, Martínez V, Pozo-Rubio T, et al. Methods to develop an in silico clinical trial: computational head-to-head comparison of lisdexamfetamine and methylphenidate. *Front Psychiatry*. 2021;12:741170. doi:10.3389/fpsy.2021.741170.
11. Lopez JP, Fiori LM, Cruceanu C, et al. MicroRNAs 146a/b-5 and 425-3p and 24-3p are markers of antidepressant response and regulate MAPK/Wnt-system genes. *Nat Commun*. 2017;8:15497. doi:10.1038/ncomms15497.
12. Menard C, Pfau ML, Hodes GE, et al. Social stress induces neurovascular pathology promoting depression. *Nat Neurosci*. 2017;20:1752–1760. doi:10.1038/s41593-017-0010-3.
13. Moingeon P, Chenel M, Rousseau C, Voisin E, Guedj M. Virtual patients, digital twins and causal disease models: paving the ground for in silico clinical trials. *Drug Discov Today*. 2023;28:103605. doi:10.1016/j.drudis.2023.103605.
14. Rush AJ, Trivedi MH, Wisniewski SR, et al. Acute and longer-term outcomes in depressed outpatients requiring one or several treatment steps: a STAR*D report. *Am J Psychiatry*. 2006;163:1905–1917. doi:10.1176/ajp.2006.163.11.1905.
15. Scheepstra KWF, Mizze MR, van Scheppingen J, et al. Microglia transcriptional profiling in major depressive disorder shows inhibition of cortical gray matter microglia. *Biol Psychiatry*. 2023;94:619–629. doi:10.1016/j.biopsych.2023.04.020.
16. Shin J, Oh H, Park JH, et al. Dysregulation of FGFR1 signaling in the hippocampus facilitates depressive disorder. *Exp Mol Med*. 2025;57:1818–1836. doi:10.1038/s12276-025-01519-9.
17. Subramanian A, Tamayo P, Mootha VK, et al. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proc Natl Acad Sci USA*. 2005;102:15545–15550. doi:10.1073/pnas.0506580102.
18. Su P, Yan S, Yang J, et al. Serum amyloid P component modulates antidepressant effects through promoting membrane insertion of the serotonin transporter. *Neuropsychopharmacology*. 2023;48:508–517. doi:10.1038/s41386-022-01449-4.

19. United States Food and Drug Administration. Dexcom G5 Mobile Continuous Glucose Monitoring System: FDA Executive Summary. 2016. Includes the regulatory history of the UVA/Padova type 1 diabetes simulator.
20. Viceconti M, Hunter P. The Virtual Physiological Human: ten years after. *Annu Rev Biomed Eng.* 2016;18:103–123. doi:10.1146/annurev-bioeng-110915-114742.
21. Viceconti M, Pappalardo F, Rodriguez B, Horner M, Bischoff J, Musuamba Tshinanu F. In silico trials: verification, validation and uncertainty quantification of predictive models used in regulatory evaluation of biomedical products. *Methods.* 2021;185:120–127. doi:10.1016/j.ymeth.2020.01.011.
22. Visentin R, Dalla Man C, Kovatchev B, Cobelli C. The University of Virginia/Padova type 1 diabetes simulator matches the glucose traces of a clinical trial. *Diabetes Technol Ther.* 2014;16:428–434. doi:10.1089/dia.2013.0377.
23. World Health Organization. Depressive disorder (depression): fact sheet. Updated 11 September 2026. <https://www.who.int/news-room/fact-sheets/detail/depression>.

10 Supplementary material

The supplementary spreadsheets are unchanged copies of the complete DEPRESSION10 exports used for the analyses. They provide full target-sensitivity, MoA, effector, responder and baseline tables beyond the rows reproduced in the manuscript. The raw repeated-dose concentration profiles are retained as pharmacokinetic quality control (Figure S1); they are not a principal biological endpoint.




Figure S1. Repeated-dose concentration profiles generated by the PBPK module. Curves document the upstream calculation from which the bounded patient exposure factor in Figure 2 was obtained.

Table 8 maps each unchanged supplementary spreadsheet group to its corresponding analysis in the manuscript.

Table 8. Supplementary spreadsheet inventory

File group	Content	Use in manuscript
TAG_ANALISIS_TARGETS_1–3.xlsx	Leave-one-target-out sensitivity	Figure 3 and Section 4.2
TAG_TABLA_FUENTESYSUMIDEROS_1–4.xlsx	Complete informative MoA proteins	Figures 4–6 and Section 4.3
TAG_TABLA_EFECTORES_SELECCIONADOS.xlsx	Depression effector catalogue	Methods and Figure 7
TAG_ANALISIS_TRATAMIENTO_PROTEINAS_1–4.xlsx	Complete responder protein contrasts	Figure 8 and Sections 4.6–4.7
TAG_ANALISIS_TRATAMIENTO_BIOMARCADORES_1,2,4.xlsx	Complete baseline contrasts	Section 4.8
TAG_TABLA_MOA_RESPONDEDORES_1–4.xlsx	Responder overlays on general MoA	Figures 10a–10d

Note. No baseline spreadsheet exists for CITAL40 because no protein passed the export criterion $q < 0.10$.