Article:
Knight, Martin, Hutchings, Michael, White, Piran Crawfurd Limond et al. (2022) A mechanistic model captures livestock trading, disease dynamics, and compensatory behaviour in response to control measures. Journal of Theoretical Biology. 111059. ISSN: 1943-5193

https://doi.org/10.1016/j.jtbi.2022.111059

Supplementary material

1 Iterative method for obtaining values of a_i and b_i

Obtaining values of a_i and b_i for each farm i is challenging, as the rates of partnership formation and trade are functions of time-varying farm-level stock quantities, namely supply $S_i(t)$ and demand $D_i(t)$. We use an iterative process to obtain a_i and b_i by taking initial values for both constants. Using these initial values, we simulate the system and obtain equilibrium values for k_i^* (the equilibrium average number of trade partners for farm i) and Φ_i^* (the equilibrium per unit-time number of trades for farm i). These values are compared to their expected values ($<k_i>$ for number of trade partners and $<\Phi_i>$ for number of trades) obtained as explained in Section X of the main text. New values for a_i and b_i are obtained by calculating the factor difference between simulated and expected values as follows:

$$a_i(s+1) = a_i(s)\frac{<k_i>}{k_i^*(s)}, \tag{1}$$

$$b_i(s+1) = b_i(s)\frac{<\Phi_i>}{\Phi_i(s)}, \tag{2}$$

where $a_i(s)$ and $b_i(s)$ are the values of a_i and b_i for farm i on simulation s, and $k_i^*(s)$ and $\Phi_i^*(s)$ are, respectively, the equilibrium number of trade partners and trades for farm i on simulation s. The new values of a_i and b_i are used to re-simulate the system, and the recursive formula above is used again to inform updated values of

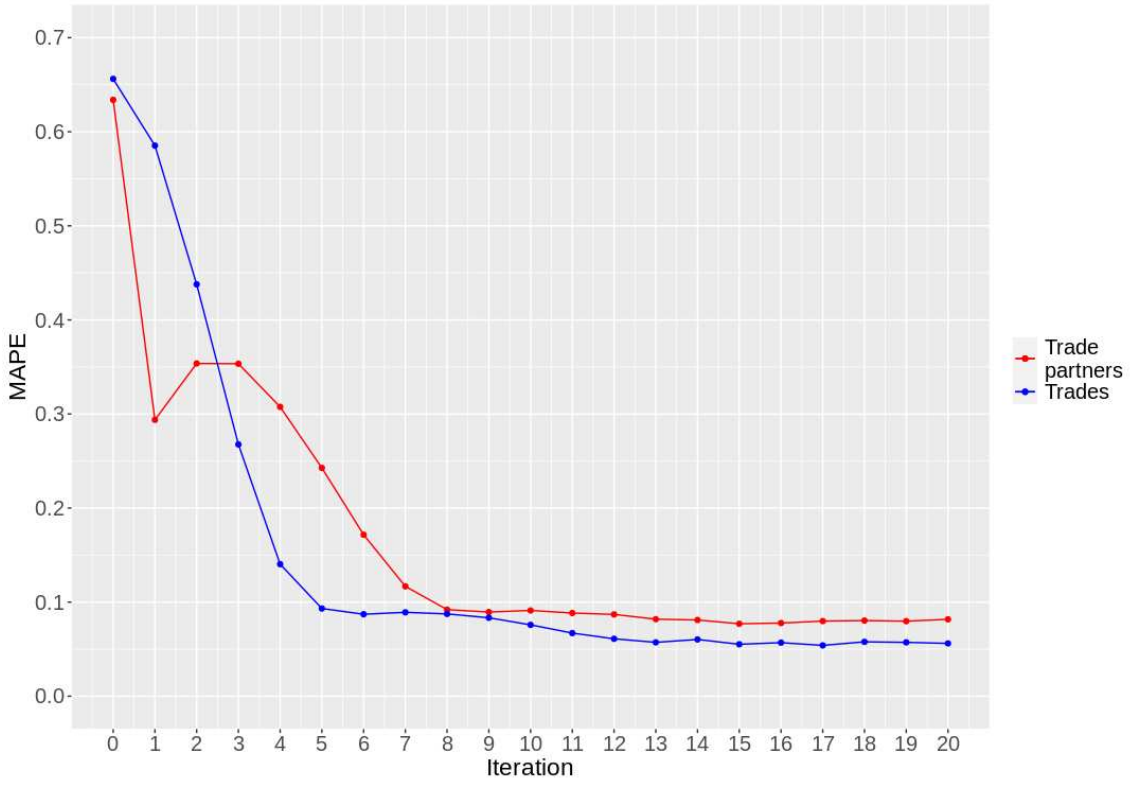

Figure 1: Mean Absolute Percentage Error (MAPE) for farms' average number of trading partners and trades comparing simulation output from a particular iteration of the parameter fitting process and the expected values drawn from distributions. Iteration 0 represents the initial state where $a_i = b_i = 0.1$ for all farms i.

a_i and b_i. This iterative process is repeated until values are obtained that accurately replicate desired trading patterns for each farm.

Our iterative method is somewhat analogous to an Expectation-Maximisation algorithm (EM algorithm). The EM algorithm calculates an expectation for the log-likelihood given current parameter estimates, followed by computation of parameter values that maximise the expectation of the calculated log-likelihood. The two steps are then iteratively repeated until suitable parameter values are obtained [1].

To assess the fits of simulation to expected values after each iteration of our fitting method, we use the Mean Absolute Percentage Error (MAPE). Figure 1 shows that our iterative method quickly finds values of a_i and b_i that minimises the error between expected and simulated values. Figure 2 further highlights this, showing after 20 iterations simulation output agrees well with desired trading patterns for farms

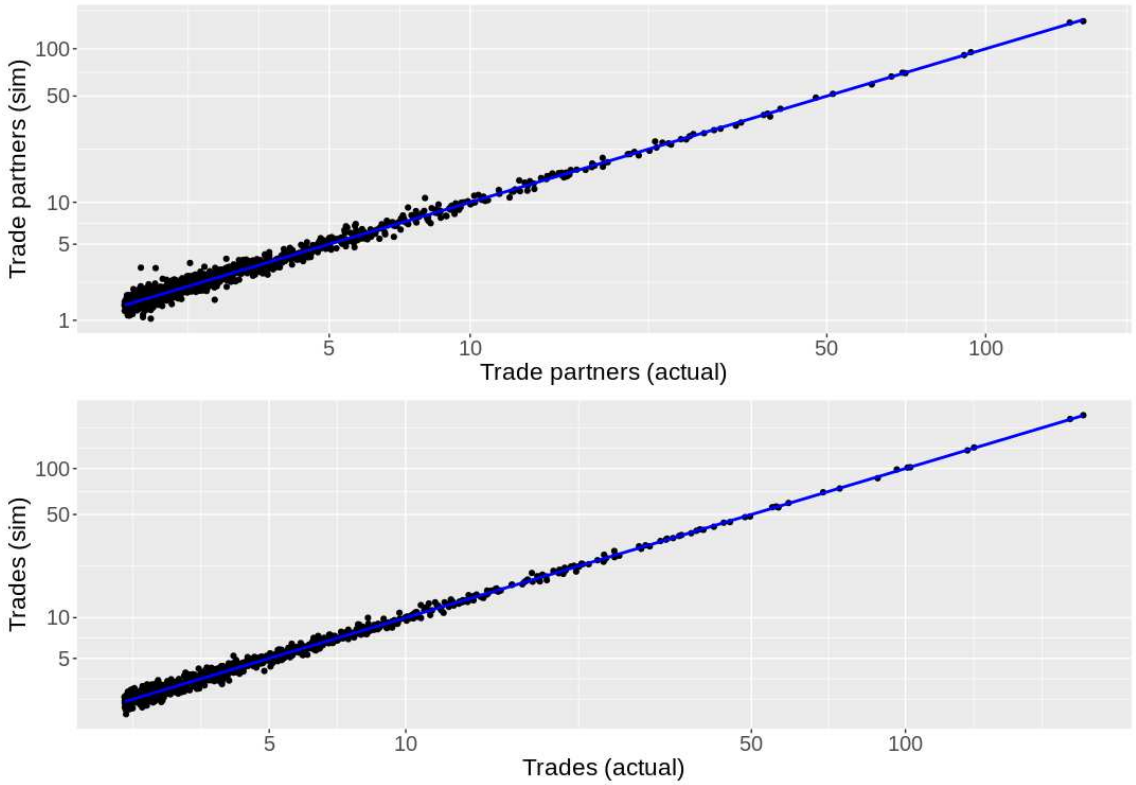

Figure 2: Per-farm fits for average number of trade partners and trades, comparing actual values (as drawn from their respective distributions) and simulated values after 20 iterations of our fitting method. Blue lines represent a linear fit. Axes are plotted on a $\log_{10}$ scale.

across the entire distribution of trading partners and number of trades.

2 Stock divergences

Figure 4 shows the dynamics of global excess demand (difference between global demand and supply) over time for a number of independent stochastic simulations in the absence of a pricing model, i.e. when the price sensitivity parameter $\sigma = 0$. In the absence of a pricing model, imbalances in supply and/or demand do not inform changes in price, and thus farm stock generations rates $\eta_i(P)$ and $\zeta_i(P)$ do not deviate from their equilibrium values η_i^* and ζ_i^*. As such, stock imbalances are not diminished, instead they accumulate over time leading to long-term divergent imbalances. Figure 4 highlights that these imbalances can occur with equal frequency for both supply and demand. Thus, our pricing model mitigates stochastic elements

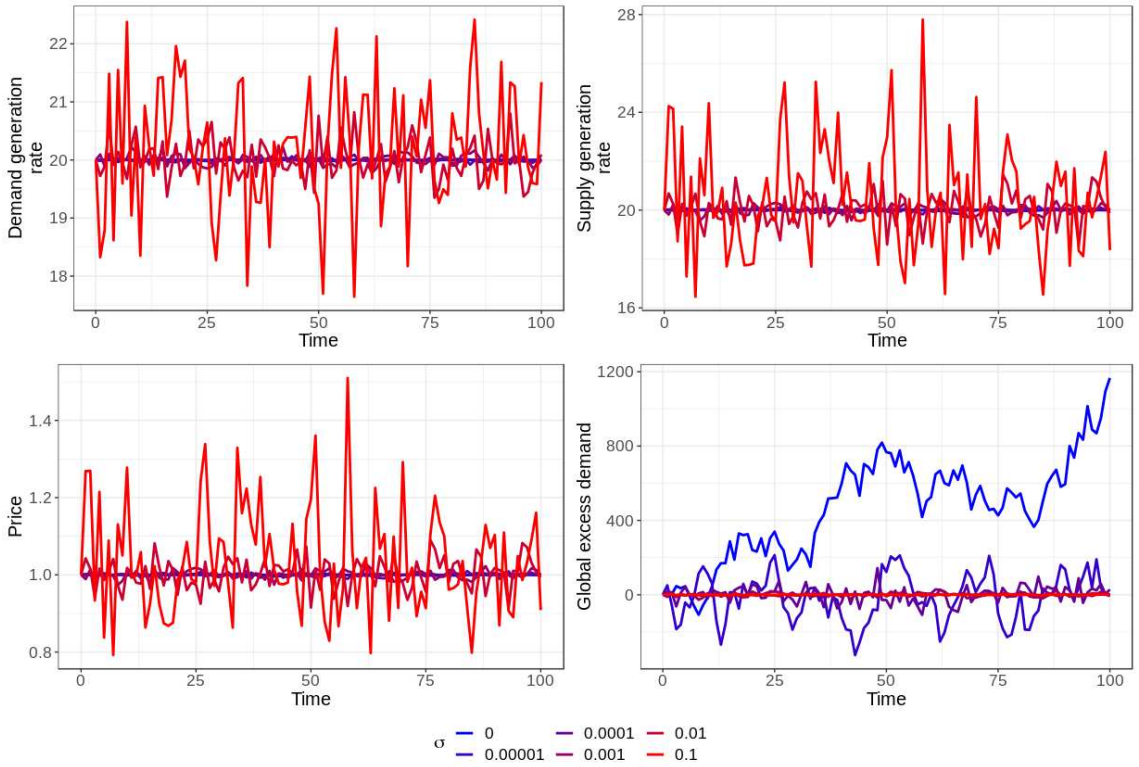

Figure 3: System-average stock generation rates, price, and global excess demand (the difference in global demand and supply) for various values of the price sensitivity parameter σ. For all values of σ, stock generation rates and price are stable, however for larger values of σ we note that price changes in response to stock imbalances are larger, which result in larger alterations to stock generation rates. Smaller values of σ result in more variable stock imbalances, and the special case when $\sigma = 0$ results in divergent stock imbalances as the system and individuals cannot alter prices and stock generation in response to these imbalances. Each trajectory is obtained by averaging over 15 independent stochastic simulations.

of our model that lead to these imbalances, ensuring supply and demand remain stable.

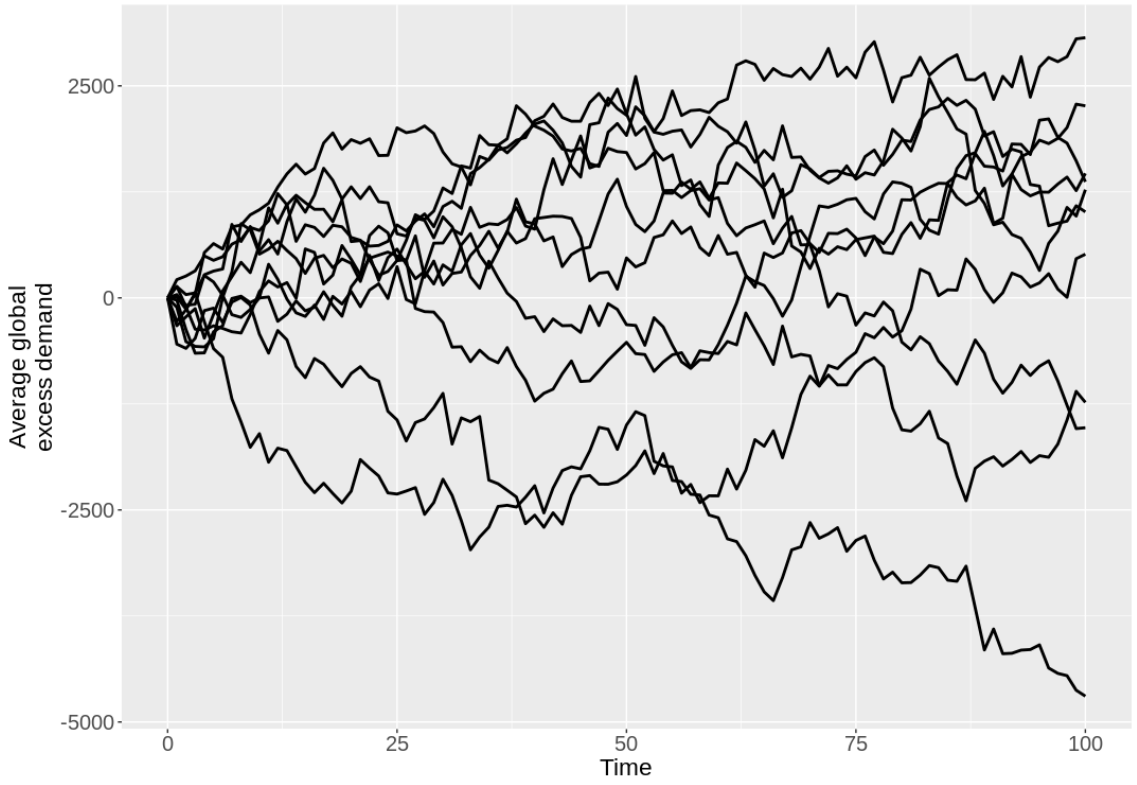

Figure 4: Global excess demand (difference between global demand and supply) for 10 independent stochastic simulations in the absence of a pricing model, i.e. when price sensitivity parameter $\sigma = 0$. Divergences are not confined to either supply or demand, occurring uniformly based on cascading imbalances that cannot be rectified by alterations in supply and demand generation rates.

3 Shocks to supply, both supply and demand, and removal of all supply and demand

3.1 Shocks to supply

Similarly to shocks to farm-level demand as described in the main text, we apply instantaneous shocks to farm-level supply at time $t = 50$. Figures 6 and 7 show the system response to these shocks. In general, shocks to supply result in increased network connectivity, as well as increased trading frequency with larger batch sizes in a manner similar to shocks in demand. These changes are temporary while the increased supply is cleared from the system, after which the system returns to pre-shock equilibrium values. In contrast with shocks to demand, shocks in supply result in increases in demand generation rate, reductions in supply generation rate,

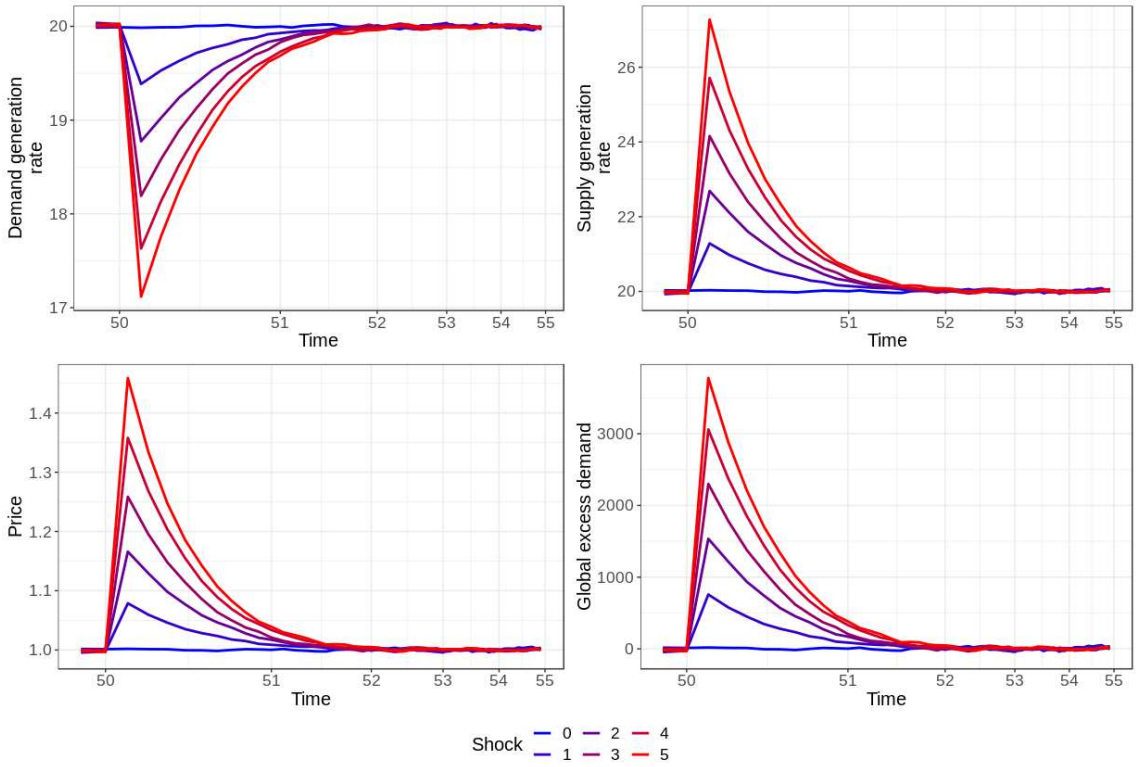

Figure 5: Stock generation rates, price, and global excess demand (difference between global demand and supply) in response to instantaneous shocks in farm-level demand. In all cases, shocks are applied at $t = 50$ and x-axes are $\log_{10}$ scaled. The size of shocks represent the magnitude by which farm-level supply is increased. Each trajectory is obtained by averaging over 15 independent simulations.

as well as reductions in price. This is to be expected given our pricing model yields reductions in price for surplus supply, which translates into increases in demand generation and reductions in supply generation.

3.2 Simultaneous shocks to supply and demand

Figures 8 and 9 shows the impact of simultaneous, and of similar magnitude, shocks to both supply and demand. In general, the trading system responds similarly to shocks in supply and demand, however the impacts on price and stock generation rates are noticeably different. In particular, there is no change when shocks are introduced, however this is to be expected as shocks are applied to supply and demand simultaneously, hence the difference is unchanged by shocks. Thus, price and stock generation rates are unchanged.

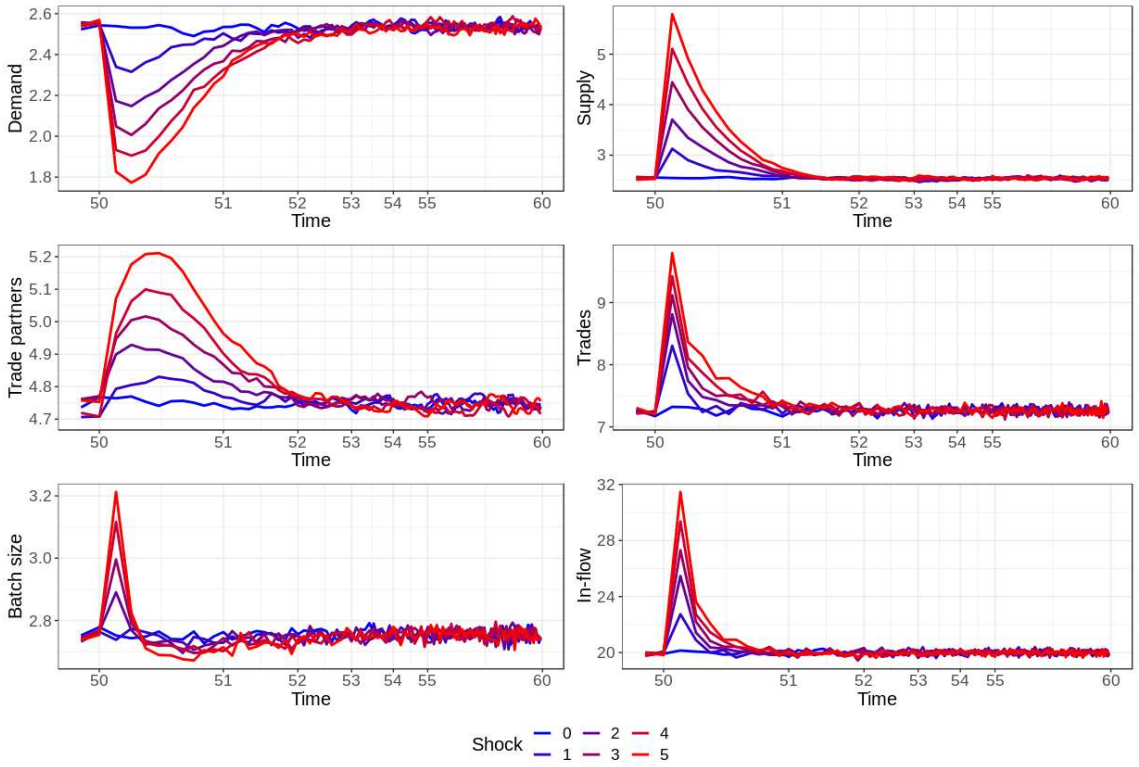

Figure 6: System-level response to instantaneous shocks in farm-level supply For a shock of size s, for all farms i supply is updated to $\mathcal{S}_i(t) \to \mathcal{S}_i(t) + s$. In all cases, shocks are applied at $t = 50$ and x-axes are $\log_{10}$ scaled. The sizes of shocks represent the magnitude by which farm-level supply is increased. Each trajectory is obtained by averaging over 15 independent simulations.

3.3 Removal of all supply and demand

Figure 10 shows the response of the trading system to the instantaneous removal of all supply and demand. As predicted, removal of all stock results in a sharp reduction in trading frequency, and the network initially becomes more disconnected. However, as farms begin to regenerate supply and demand, the system rapidly returns to the pre-shock state and there are no long-term alterations to the trading system.

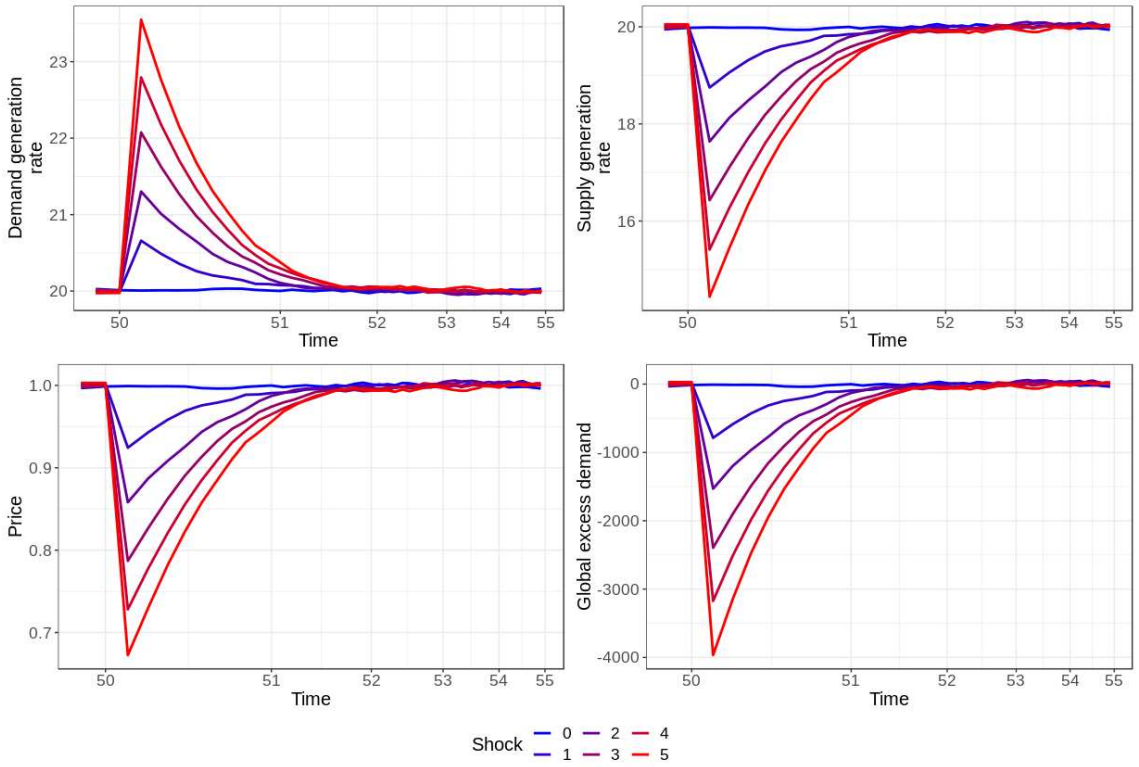

Figure 7: Stock generation rates, price, and global excess demand (difference between global demand and supply) in response to instantaneous shocks in farm-level supply. In all cases, shocks are applied at $t = 50$ and x-axes are $\log_{10}$ scaled. The size of shocks represent the magnitude by which farm-level supply is increased. Each trajectory is obtained by averaging over 15 independent simulations.

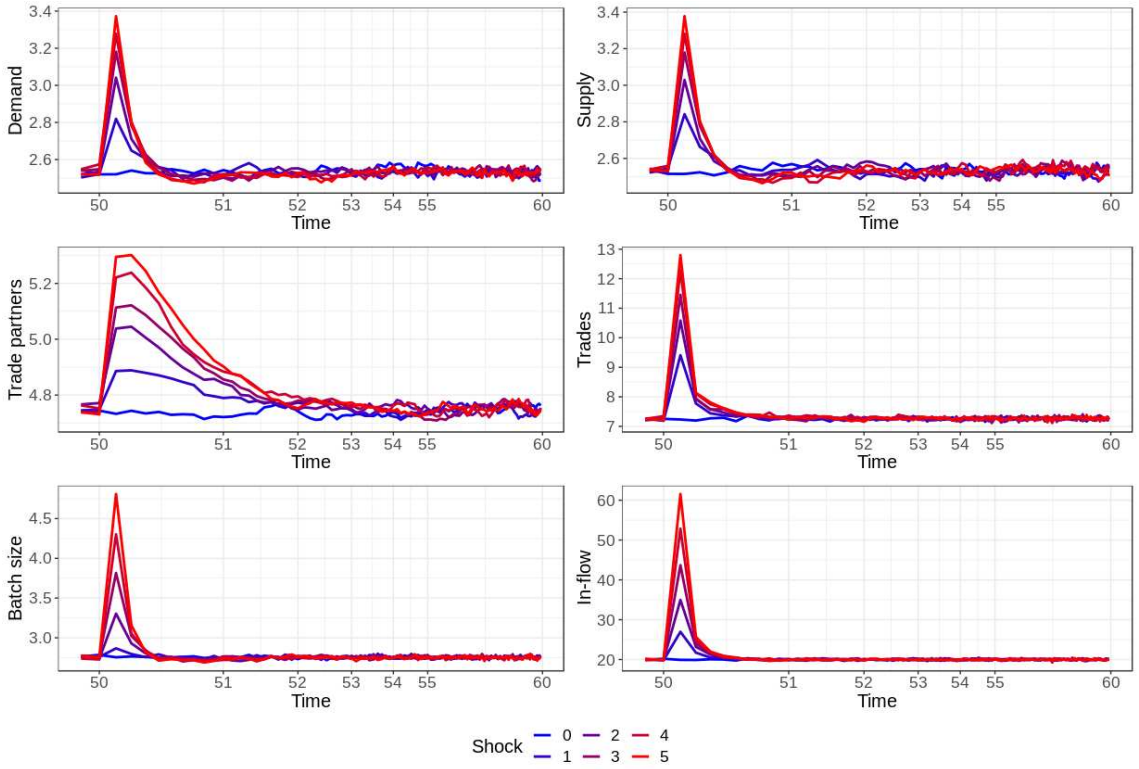

Figure 8: System-level response to instantaneous simultaneous shocks in farm-level supply and demand. For a shock of size s, for all farms i supply and demand is updated to $\mathcal{S}_i(t) \to \mathcal{S}_i(t) + s$ and $\mathcal{D}_i(t) \to \mathcal{D}_i(t) + s$, respectively. In all cases, shocks are applied at $t = 50$ and x-axes are $\log_{10}$ scaled. The sizes of shocks represent the magnitude by which farm-level stock quantities are increased. Each trajectory is obtained by averaging over 15 independent simulations.

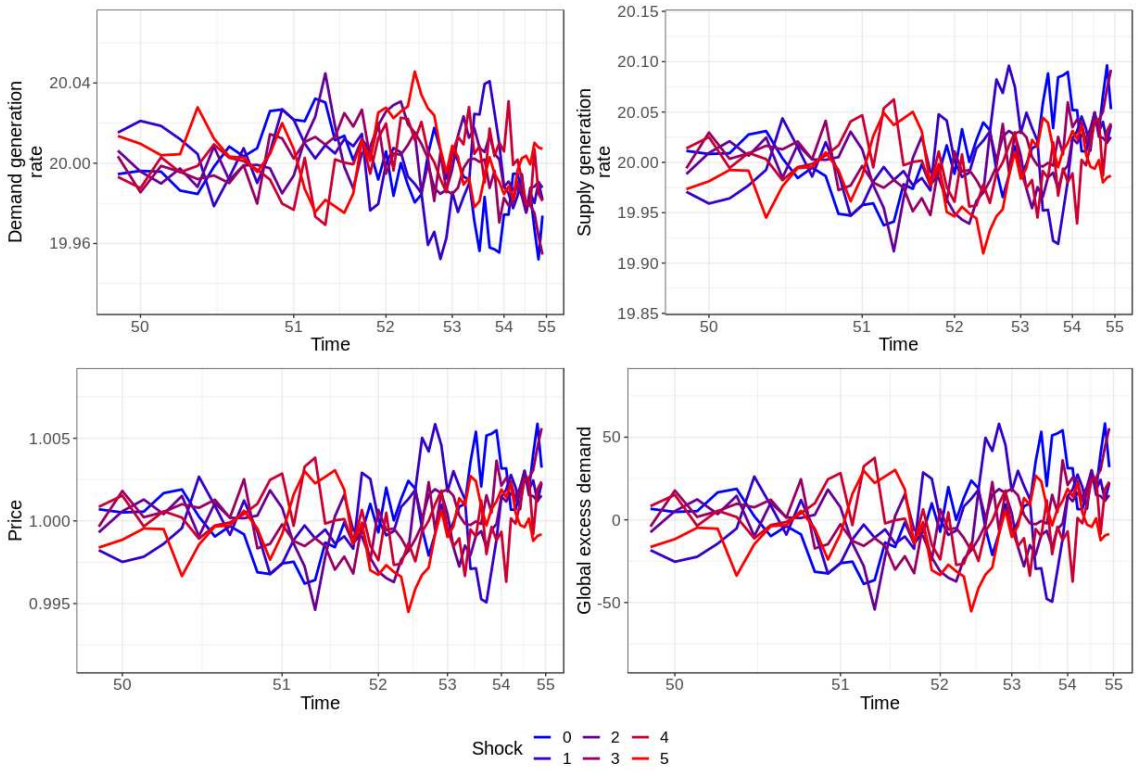

Figure 9: Stock generation rates, price, and global excess demand (difference between global demand and supply) in response to instantaneous simultaneous shocks in farm-level supply and demand. In all cases, shocks are applied at $t = 50$ and x-axes are $\log_{10}$ scaled. The size of shocks represent the magnitude by which farm-level stock quantities are increased. Each trajectory is obtained by averaging over 15 independent simulations.

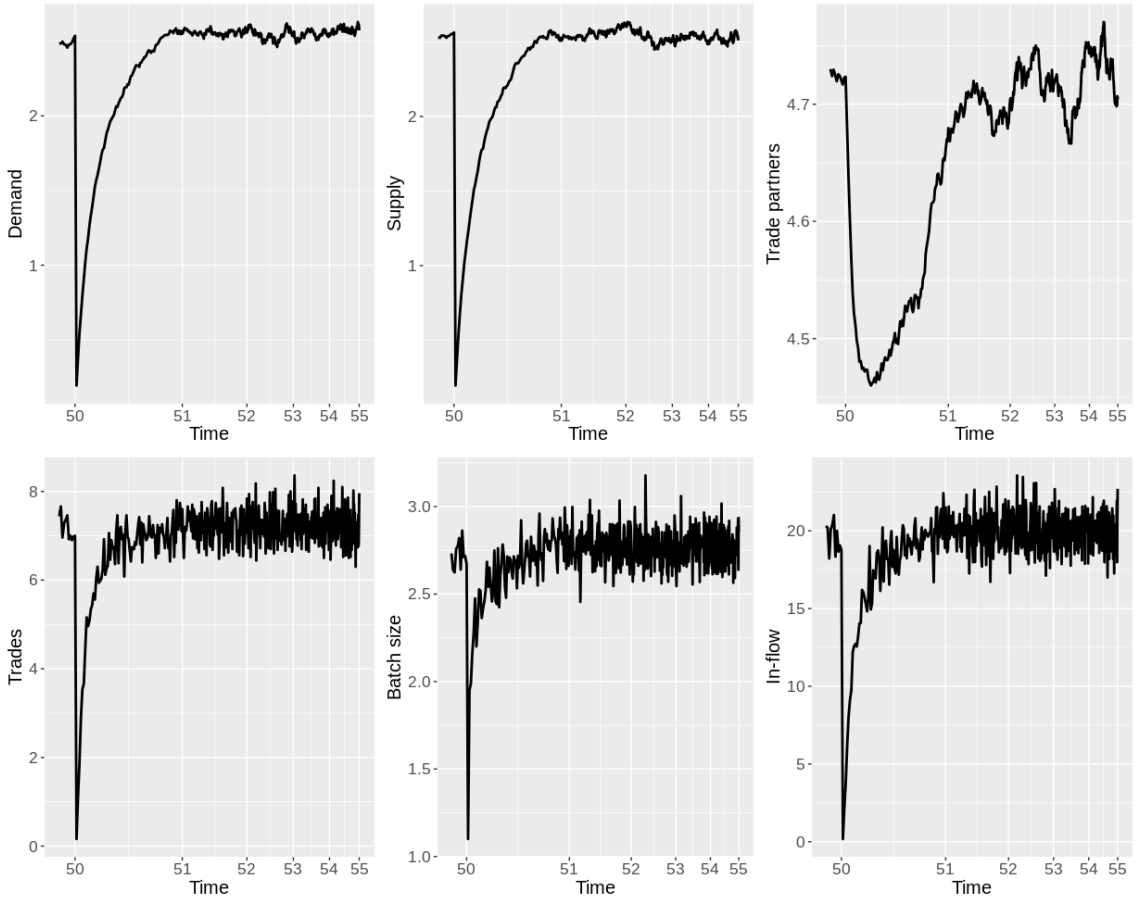

Figure 10: System-level response to instantaneous removal of all supply demand. Stock removal is applied at $t = 50$ and x-axes are $\log_{10}$ scaled. Each trajectory is obtained by averaging over 5 independent simulations.

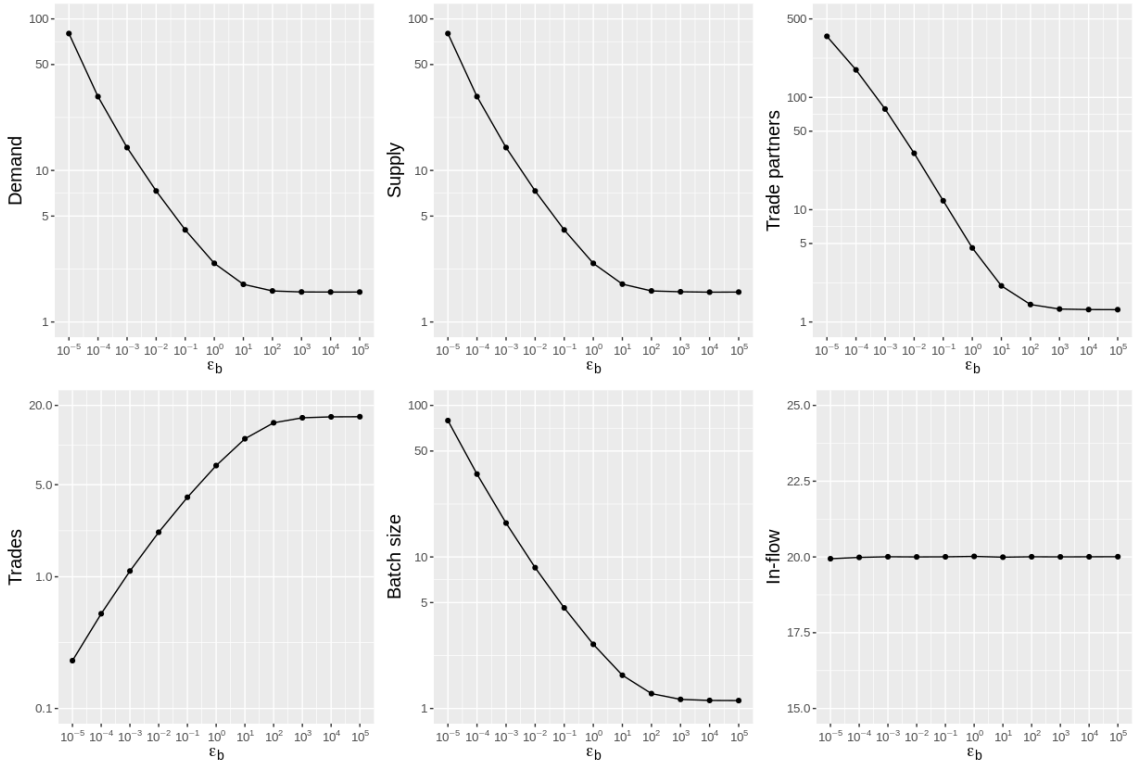

Figure 11: Equilibrium values for system-average trade quantities for various values of the trade frequency scaling parameter ε_b. Initially, the system reaches an equilibrium at baseline trading patterns when $\varepsilon_b = 1$ before changes to ε_a are made at $t = 50$. Averages are obtained by averaging from $t = 75$ to $t = 100$.

4 Effect of long-term changes to partnership formations and partnership durations on trading system

4.1 Changes to ε_a

Here we explore the effect of alterations to the propensity for farms to form new trade partnerships, via the scaling parameter ε_a. As described in the main text, this parameter linearly scales the per-farm constant a_i such that small ε_a decreases the propensity for farms to form trade partnerships, and vice versa.

Figure 12 shows the impact on the trading system of various values of ε_a over time, and Figure 13 shows the long-term equilibrium values for the trading system once

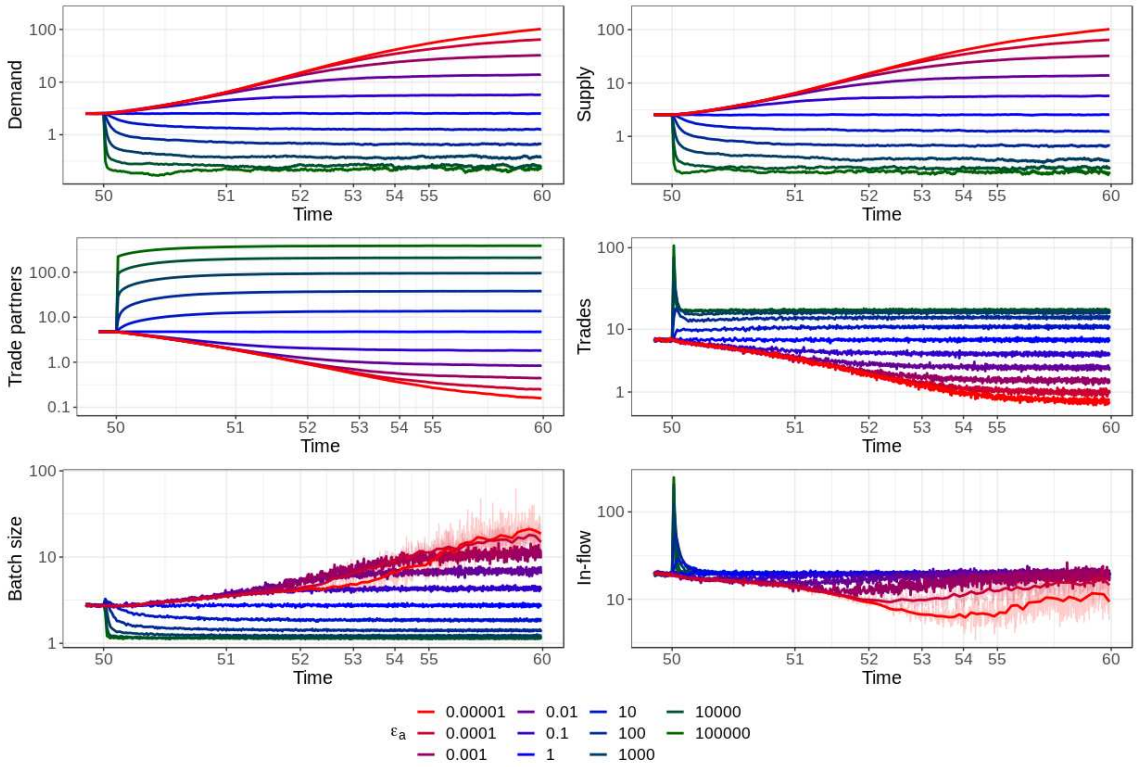

Figure 12: System trading quantities for various values of the scaling factor to the partnership formation rate constant a_i, ε_a. In all cases, the system is allowed to reach an equilibrium at $\varepsilon_a = 1$ before changes to ε_a are made at $t = 50$. For $\varepsilon_a = 0.0001$ and $\varepsilon_a = 0.00001$, faded lines are simulation output and solid lines are fits using a local polynomial regression. Axes are $\log_{10}$ scaled.

the system has adapted to the changes in partnership formation imposed by various values of ε_a. We initially allow the system to reach an equilibrium at baseline trading patterns, i.e. when $\varepsilon_a = 1$, before changes are instantaneously applied at $t = 50$. For small ε_a, we notice initial reductions in the average number of trade partners farms have, and supply and demand begin to increase. We note that the reduction in network connectivity is gradual for $\varepsilon_a < 1$ because this is determined by the partnership cessation rate, δ_i for farm i, which is a constant unaffected by other trading factors such as stock levels. As supply and demand levels continue to increase, this pressure results in network connectivity and trading frequency to stabilise at new equilibrium values, and supply and demand also begin to reach a new equilibrium, albeit farms have greater available supply and unmet demand compared with the baseline case $\varepsilon_a = 1$. This results in larger batch sizes, and this interplay allows for farm flows to return to baseline equilibrium values following

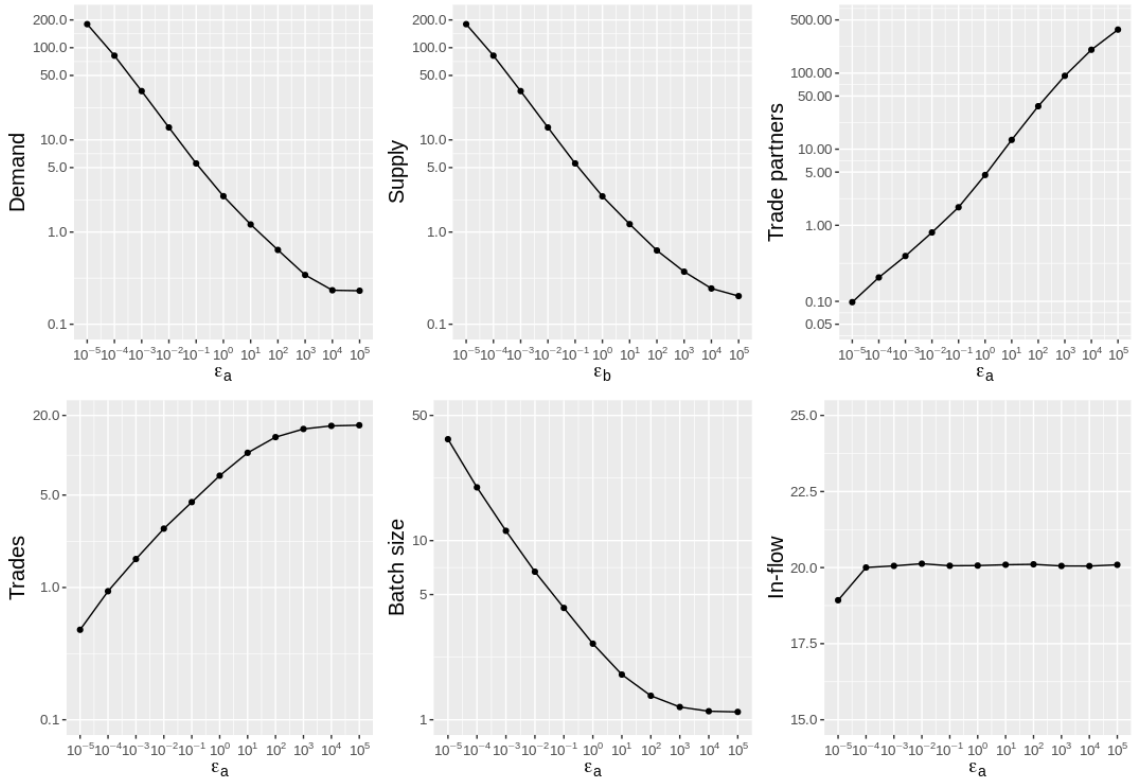

Figure 13: Equilibrium values for system-average trade quantities for various values of the partnership formation scaling parameter ε_a. Initially, the system reaches an equilibrium at baseline trading patterns when $\varepsilon_a = 1$ before changes to ε_a are made at $t = 50$. Averages are obtained by averaging from $t = 75$ to $t = 100$.

an initial drop. The system takes longer to adapt to small ε_a than large ε_a. In these cases, network connectivity rapidly increases and finds a new equilibrium, as does trading frequency. This allows unmet demand to rapidly decrease and farms similarly have smaller available supply, resulting in smaller batch sizes. Despite an initial increase in farm flows, this clear-out of supply and demand quickly brings farm flows back to baseline equilibrium values.

From Figure 13 we notice similar exponential relationships between ε_a and various trading quantities as was observed for ε_b. In addition, this exponential relationship appears to be present only for $\varepsilon_a < 10^2$, due to the biological constraints on the batch size taking minimum size 1 imposed by our model being reached for values of ε_a greater than this.

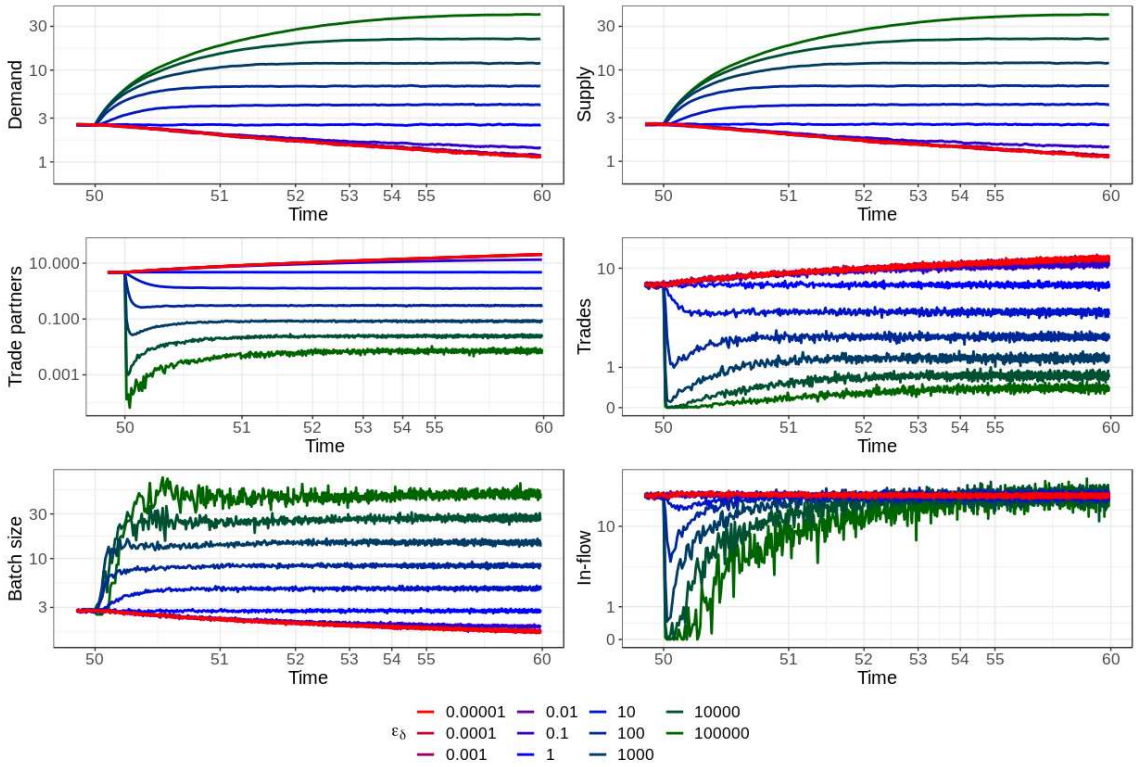

Figure 14: System trading quantities for various values of the scaling factor to the partnership cessation rate constant δ_i, ε_δ. In all cases, the system is allowed to reach an equilibrium at $\varepsilon_\delta = 1$ before changes to ε_δ are made at $t = 50$. Axes are $\log_{10}$ scaled.

4.2 Changes to ε_δ

We now explore the effect of altering the average partnership duration via ε_δ. In this case, small ε_δ results in increases in partnership durations, and vice versa.

Figures 14 shows the impact on the trading system of various values of ε_δ over time, and 15 shows the long-term equilibrium for the trading system once the system has adapted to the changes in partnership durations imposed by various values of ε_δ. As with ε_b and ε_a, we initially allow the system to reach equilibrium at $\varepsilon_\delta = 1$ before changes are made. For small ε_δ, trade partnerships durations are increased, with network connectivity increasing. This results in increased trade and supply and demand begin to decrease (batch sizes decrease as a result). Conversely, for large ε_δ trade partnerships become increasingly short lasting. As such, there is an initial rapid reduction in network connectivity and trading frequency, and supply and demand begin to increase. However, the increase in stock levels result in system

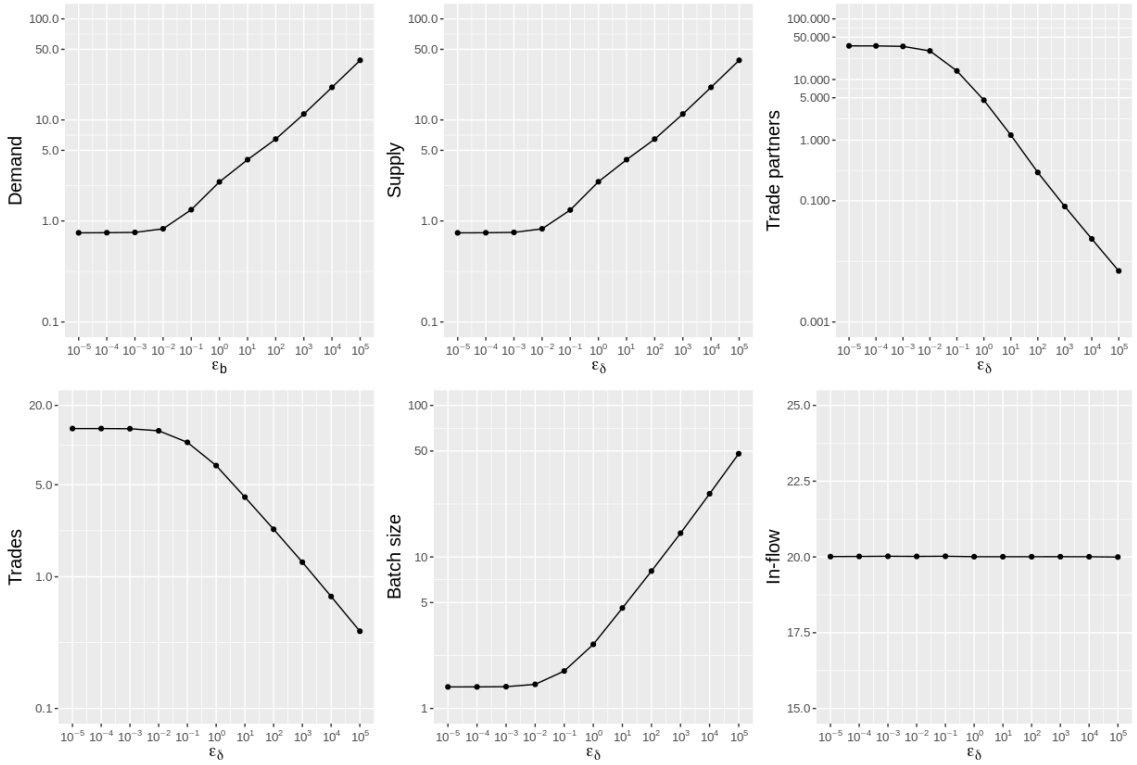

Figure 15: Equilibrium values for system-average trade quantities for various values of the partnership cessation scaling parameter ε_δ. Initially, the system reaches an equilibrium at baseline trading patterns when $\varepsilon_\delta = 1$ before changes to ε_δ are made at $t = 50$. Averages are obtained by averaging from $t = 75$ to $t = 100$.

adaptation and network connectivity stabilises to a new equilibrium, and trading frequency begins to increase and also finds a new equilibrium. Stock levels also find a new equilibrium, however farms have increased available supply and unmet demand. Farm flows exhibit similar behaviour to network connectivity and trading frequency, initially decreasing followed by a rebound and return to the original equilibrium at baseline trading patterns (when $\varepsilon_\delta = 1$).

Figure 15 shows, as before, an exponential relationship between ε_δ and various trading quantities. Similarly, we observe that this exponential relationship does not hold for all ε_δ considered here. For small ε_δ the biological constraint of small batch sizes is increasingly encountered, so that trading patterns do not alter significantly as ε_δ is further decreased (at least in the long-term).

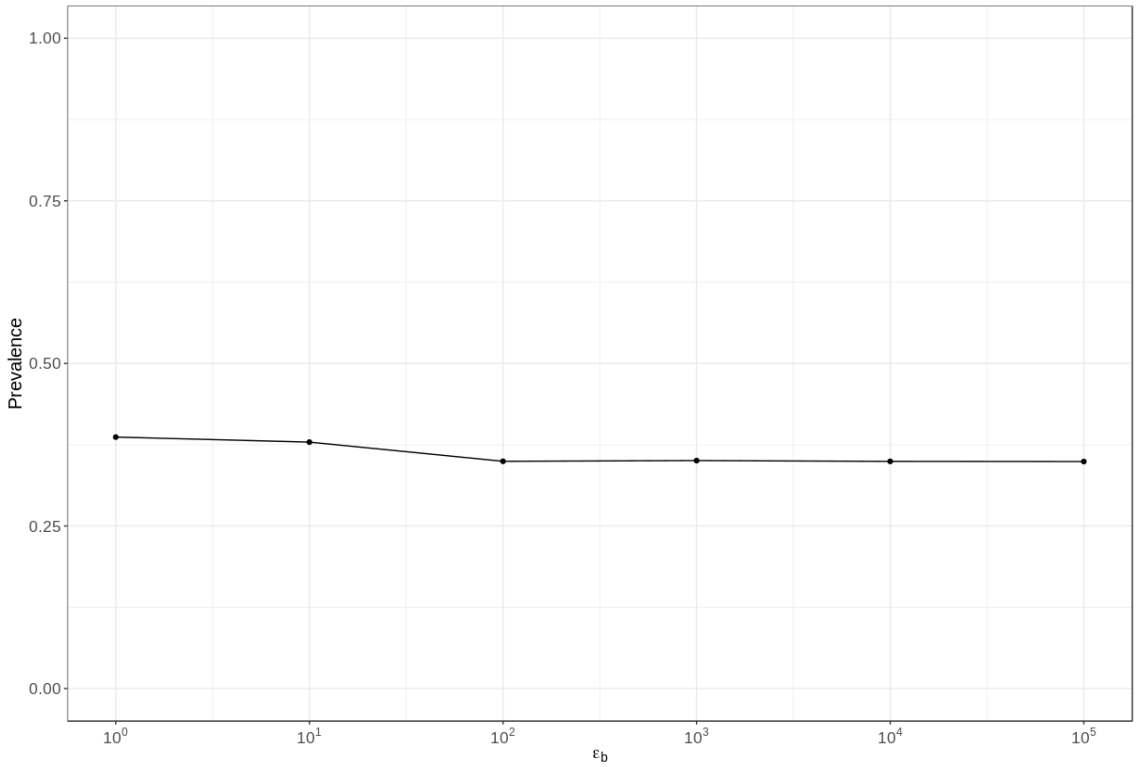

Figure 16: Equilibrium prevalence for various values of the rate scaling parameter ε_b for disease parameters $\lambda = 0.1$ and $\gamma = 2/3$. Changes to ε_b are made at $t = 50$ after disease is introduced and allowed to reach an equilibrium at baseline trading patterns, i.e. when $\varepsilon_b = 1$. X-axis is log_{10} scaled.

5 Effect of long-terms changes to partnership formations and partnership durations on prevalence

5.1 Changes to ε_a

We here explore the effect of changes to the propensity for farms to form new trade partnerships on long-term disease prevalence. This is done by scaling the partnership formation rate by a constant ε_a in a similar manner to the trade rate as described in the main text. Figure 17 shows that farms must essentially cease forming new trade partnerships to completely remove disease, and even to reduce prevalence by significant magnitudes. As with the trade rate, and as described in Section 4.1 here, this arises as a consequence of individual adaptation to find new avenues of trade

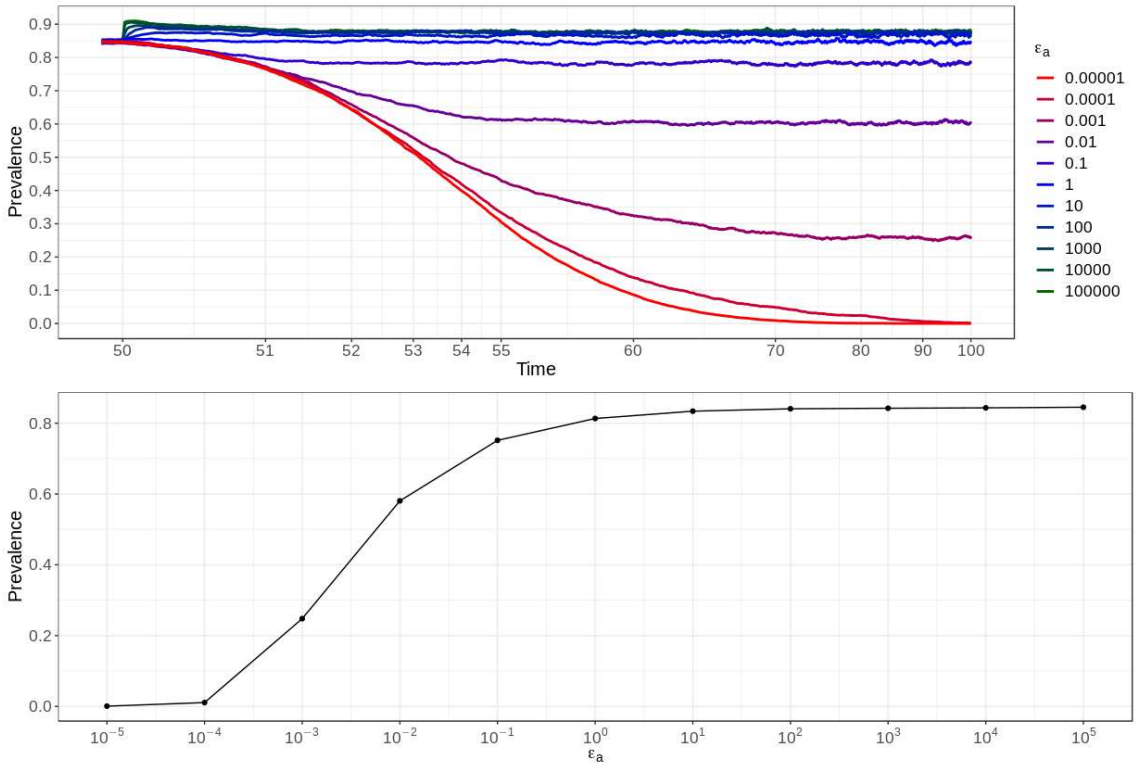

Figure 17: Disease prevalence over time (top panel) and equilibrium prevalence (bottom panel) for various values of the partnership formation scaling parameter ε_a. Changes to ε_a are made at $t = 50$ after disease is introduced and allowed to reach an equilibrium at baseline trading patterns, i.e. when $\varepsilon_a = 1$. For top panel, x-axis is $\log_{10}$ scaled, and for bottom panel both axes are $\log_{10}$ scaled.

that maintain animal flows. Increasing ε_a increases the connectivity of the network, however leads to only small increases in prevalence that do not continuously scale with ε_a.

5.2 Changes to ε_δ

Alterations to the average duration of trade partnerships, made by scaling the partnership cessation rate δ_i for farm i by a constant ε_δ. Figure 18 shows that increasing the duration of trade partnerships does reduce disease prevalence, however even in the extreme cases in which partnerships are effectively permanent, disease cannot be fully eradicated. Moreover, reducing the duration of partnerships only slightly increases disease prevalence, but does not scale continuously with ε_δ.

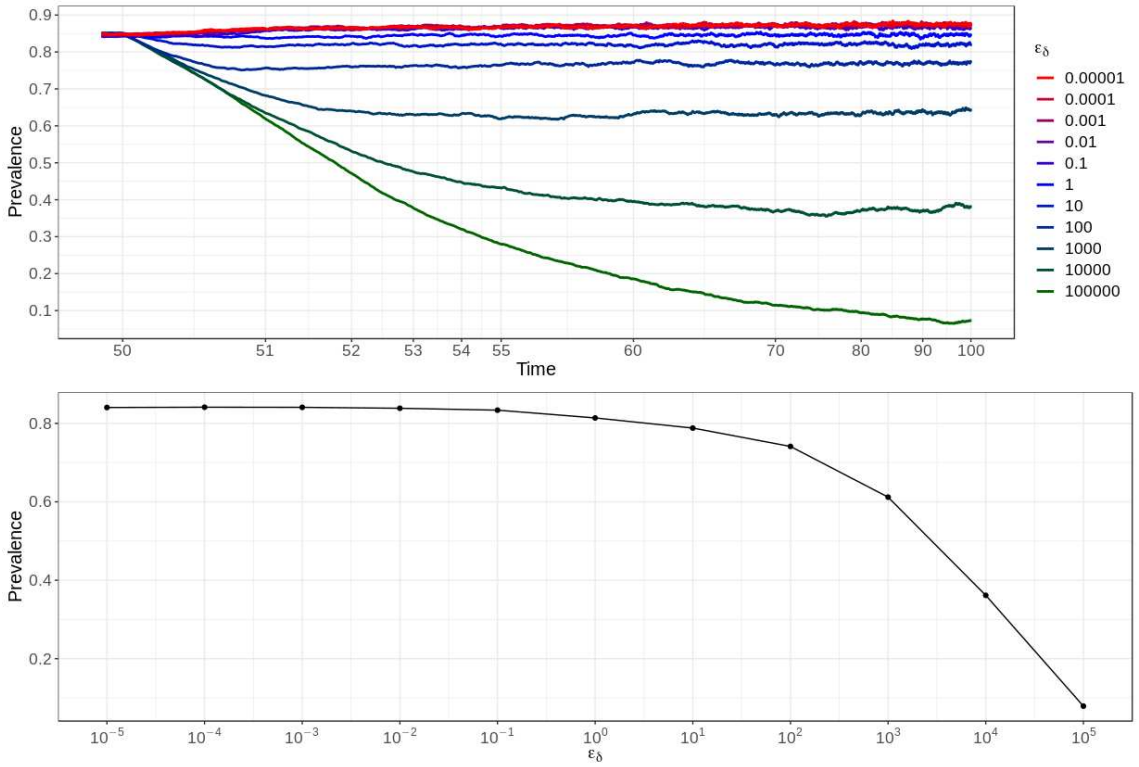

Figure 18: Disease prevalence over time (top panel) and equilibrium prevalence (bottom panel) for various values of the partnership duration scaling parameter ε_δ. Changes to ε_δ are made at $t = 50$ after disease is introduced and allowed to reach an equilibrium at baseline trading patterns, i.e. when $\varepsilon_\delta = 1$. For top panel, x-axis is $\log_{10}$ scaled, and for bottom panel both axes are $\log_{10}$ scaled.

6 Investigating impact of individual adaptation on disease prevalence

Here we show the effect of adaptation in response to changes in ε_b on between-herd disease prevalence by comparing with our previous trading model [2]. For this model, farms are not characterised by at time t stock levels, nor do we include a pricing model. Instead, farms' rates of partnership formation and trade are functions of long-term expected animal in-flows, η_i^* and ζ_i^* from the main text. The rates are

$$\alpha_{ij} = \frac{a_i}{N}\eta_i^* \left(\zeta_j^*\right)^m,$$
$$\varphi_{ij} = b_i \min(\eta_i^*, \zeta_j^*)$$

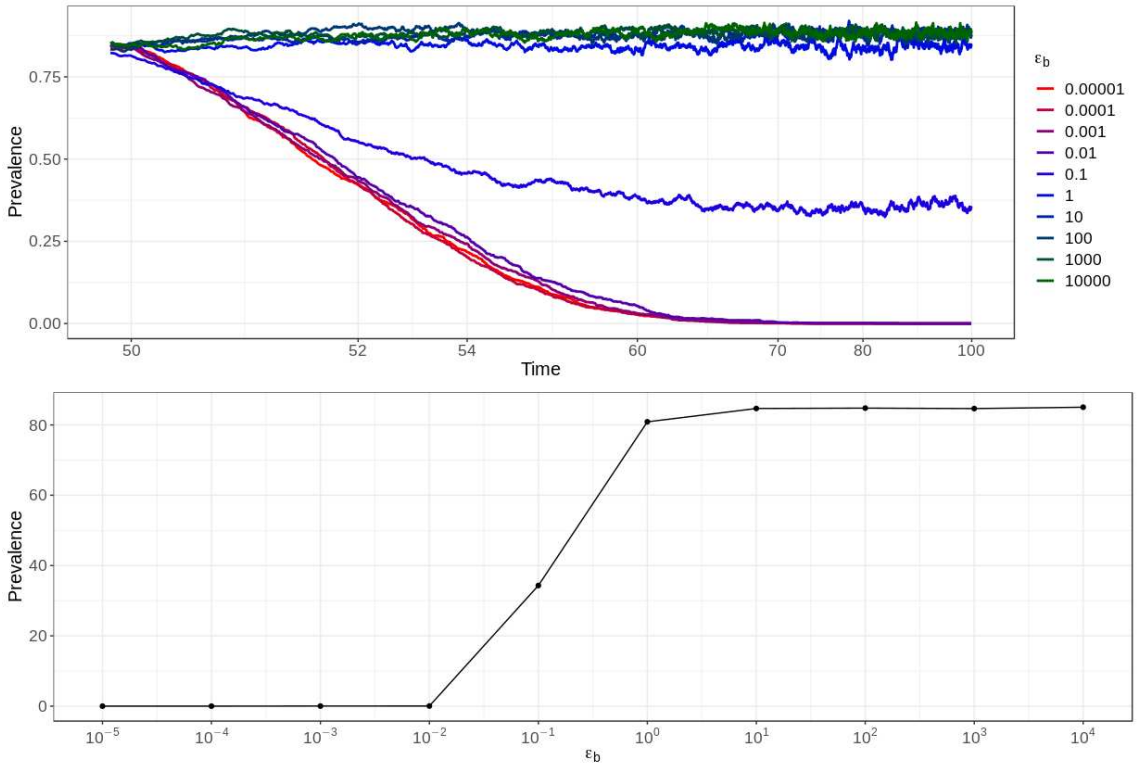

Figure 19: Between-herd disease prevalence for various values of ε_b under the trading model of [2]

for the partnership formation rate and rate of trade, respectively. Batch sizes θ_i are also constant and set by the buying farm. As such, regardless of the selling farm, farm i will always purchase θ_i animals during a trade. To ensure a meaningful comparison between the two models, we maintain the same distributions as described in the main text so that a farm i in the above model maintains the same expected number of trading partners and per unit-time trades. As such, a_i and b_i are scaled appropriately so that this condition is met. We note that to ensure animal flows are maintained in the above model, we assume any scaling to b_i by ε_b is accompanied by an inversely proportional change to the batch size θ_i. At baseline trading patterns, i.e. when $\varepsilon_b = 1$, both models predict similar between-herd disease prevalence.

Figure 19 shows that under the above model, reductions in b_i via ε_b result in greater reductions in disease prevalence, and can eradicate disease for a greater range of ε_b, than was observed for similar reductions as presented in the main text. Moreover, increases in ε_b increase disease prevalence by a larger magnitude than presented in the main text. These large differences are ascribed to the restructuring and adapta-

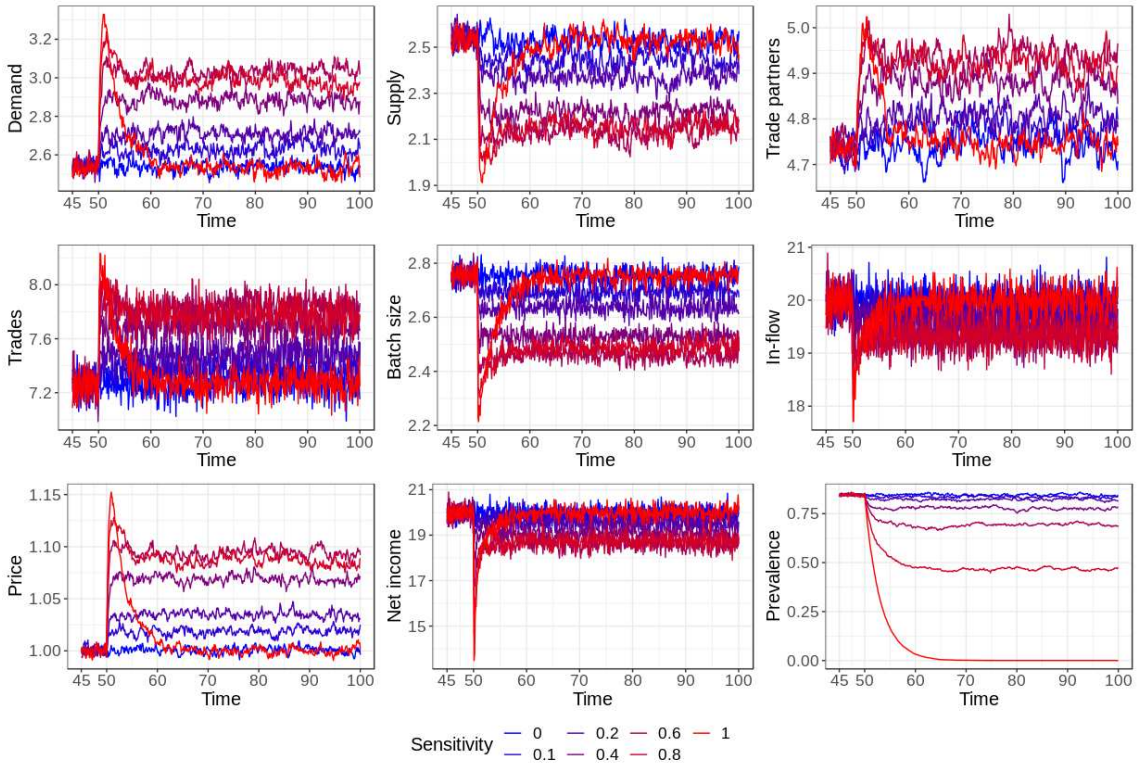

Figure 20: System average trade quantities for individual animal rejection under various test sensitivities. Each trajectory is obtained by averaging over 15 independent stochastic simulations.

tion of the trading system that occurs for the model in the main text, with changes in network connectivity as well as increasing batch sizes observed. In contrast, for the model described here, farms adapt to reductions in trading frequency by simply increasing batch sizes and there are no changes to the connectivity of the network.

7 Post-movement testing and individual animal removal

As described in the main text, we here explore the impact of post-movement testing in which individual animals in a traded batch are removed if they are detected for infection. This is in contrast to the whole batch removal in which the entire traded batch is removed if a single animal tests positive for infection.

Figure 20 shows the impact on the trading system and disease prevalence caused by the introduction of individual animal rejection. In general, the impact on the system

is smaller compared with whole batch rejection as individual animals are rejected rather than whole batches. This results in smaller disturbances to price, and as such the disturbance to stock generation rates are also smaller. However, individual animal rejection is less effective at reducing disease prevalence compared with whole batch rejection, requiring near 100% test sensitivity to fully remove disease from the system. As with whole batch rejection, we observe long-term disturbances to the trading system if disease cannot be fully removed, and temporary disturbances if disease can be fully, or near, eradicated.

References

[1] A. P. Dempster, N. M. Laird, and D. B. Rubin. "Maximum Likelihood from Incomplete Data via the EM Algorithm". In: *Journal of the Royal Statistical Society. Series B (Methodological)* 39.1 (1977), pp. 1–38.

[2] Martin A. Knight et al. "Generative models of network dynamics provide insight into the effects of trade on endemic livestock disease". In: *Royal Society Open Science* 8.3 (2021), p. 201715. DOI: 10.1098/rsos.201715.