Appendix I.

An interesting prediction of the old model is that there is non-monotonicity in the signal discriminatory capabilities of the network as a function of the ZAP70 binding kinetics **(Figure S3/S4)**. In other words, if we fix a binding rate for ZAP70 to phosphorylated ITAM and vary the unbinding rate, we will see the signal discrimination increase for intermediate values and then decrease when the off rate becomes too large. To calculate signal discriminatory capability of the null model over a wide range of ZAP70 binding kinetics, we use as a metric of discriminatory capability the area under the ROC curve[6] (Receiver Operator Characteristic) based on plots of (probability of detection, probability of false alarm) for various choices of a level of phosphorylated ZAP70 which sets the threshold for activation. We can think of this as a plot of (PD,PF) for every possible deterministic decision rule a T-cell can make, where a decision rule is uniquely defined by a threshold of activated ZAP70 that must be reached before an immune response is mounted.

This result at first glance may seem strange since the McKeithan model of kinetic proofreading suggests that faster off rates should improve discriminatory abilities. However, the important point is that the McKeithan model actually predicts that a larger *difference* in off rate between self and non-self improves discrimination. By increasing the off rate of ZAP70, we are effectively making it harder for T cells engaged with self and non-self peptides alike to reach a state of activated ZAP70.

An intuitive explanation for this result is that when the ZAP70 off rate is too low, as soon as it binds it is likely to remain bound long enough to get phosphorylated regardless of whether or not the T-cell is exposed to self or non-self antigen. Thus, activated ZAP70 concentrations in the presence of self or non-self may be similar. If the ZAP70 off rate is too high, neither self nor non-self peptide can result in ZAP70 activation, so again the activated ZAP70 concentrations in both cases are the same (very low). Thus, the model predicts an intermediate regime of ZAP70 kinetics that maximizes signal discrimination **(FIGURE S3).** We note that this simple and intuitive result can be captured by a discrete-space Markov chain.

To explain the non-monotonicity in the signal discrimination capability of the old model as a function of ZAP's unbinding rate, we built a very simple four state continuous time Markov chain to represent an individual TCR. The four states include a TCR unbound to pMHC (state 1) which transitions to TCR-pMHC complex (state 2). When the TCR-pMHC complex has formed, Zap70 can bind (state 3). It is possible for pMHC to unbind after ZAP70 has bound, leading to state 4, which has Zap70 bound to a TCR with no pMHC present; ZAP70 can also unbind from the TCR-pMHC complex before the pMHC does, leading back to state 2. We use state 3 as a proxy for activated T-cell in this simplified picture, since in the actual T-Cell, ZAP70 must be bound simultaneously with Lck associated coreceptor for activation.

We calculate two different rate matrices for self and non-self by using the same parameters for all transitions except for pMHC off rate, which is chosen to be larger in the case of non-self. We then calculate the steady state distribution by finding the kernel of the rate matrix. The steady state distributions under self and non-self allow us to compute P(activation|self) and P(activation|non-self), where P(activation) is equivalent to the steady state probability of the markov chain

occupying state 3. A plot of the likelihood ratio as a function of ZAP70's off-rate shows non-monotonicity for a fixed ratio of Koff(non-self)/Koff(self) and at a fixed Kon(Zap70) **(Figure S1).**

To ensure that the non-monotonicity is seen when considering metrics besides just the likelihood ratio, we also calculated the mutual information between the random variables S the signal={self,non-self}, and R the response={activated or not activated}. Calculating the mutual information requires knowledge of the prior probabilities over S={self,non-self}, but the channel capacity can be calculating by maximizing mutual information over the prior probability. The channel capacity is zero when response is independent of signal (i.e. when both self and non-self lead to equal odds of activation), and is maximal when the signal is a deterministic function of the input, i.e. when p(activation|self)=0, p(activation|non-self)=1 (in which case the prior probability that maximizes the information is $50/50$ over self and non-self). The channel capacity can be derived by writing down mutual information as a function of p, the prior probability for encounters with non-self peptide, and maximizing. The p* maximizer is plugged back in to the expression for mutual information to give:

$$I(p^*) = \frac{h_1 e_0 - h_0(1-e_1)}{(1-e_0-e_1)} + \log\left(1 + 2^{(h_0-h_1)/(1-e_0-e_1)}\right)$$

where e_0=P(not-activated|non-self), e_1=P(activated|self) (these quantities are derived from the steady states of the Markov chain) , and h_0,h_1 are entropies of a Bernoulli random variable parametrized by e_0,e_1:

$$h_i = -e_i \log(e_i) - (1-e_i)\log(1-e_i)$$

The channel capacity is non-monotonic as a function of ZAP70's kinetic off rate at a fixed on rate **(Figure S2).**

In the text, we calculated the AUROC for the null model based on plots of detection probability and false alarm probability for various thresholds. To quantify signal discrimination, we must choose a metric that is agnostic to the particular prior probabilities of encountering self and non-self, since such priors (though undoubtedly skewed towards encounters with self) vary by individual and are most likely not stationary. Almost every metric of binary signal discrimination used in decision theory takes into account both a probability of detection and a probability of false alarm. The probability of detection in our model is the probability for a T-cell to mount a response in the presence of foreign antigen, while the probability of false alarm is the probability that a T-cell incorrectly activates in response to self pMHC. A decision rule could be deterministic, in which case it takes the form of a sharp threshold above which the T-cell decides to mount a response (in our case the threshold represents a sufficiently high concentration of phosphorylated ZAP70), or stochastic, meaning that the T-cell does not always respond in the same way to the same concentration of phosphorylated ZAP70, but responds in some instances and not in others.

Determining the optimal decision rule would require detailed knowledge of the prior probability distribution over self/non-self and a cost function that specifies the cost of various types of mistakes (namely, not detecting foreign antigen, and incorrectly responding to self). In this case, the optimal decision rule

on the part of a T cell will be of the form of a likelihood ratio, and implies that the decision to mount a response occurs only if the amount of activated ZAP exceeds a threshold that is a function of the prior probabilities and cost function. This leads to the Bayes Least Cost decision rule[4]. (In certain cases in which there are outcomes that set the likelihood ratio equal to the optimal threshold, a randomized decision rule is optimal). If one does not know the prior probabilities but is instead willing to maximize the detection probability while insisting the false alarm probability remain below a certain threshold, the Neyman-Pearson lemma states that the optimal decision rule will still be of the form of a likelihood ratio with its own associated threshold[5].

In our case, it is not clear what the prior probabilities should be nor do we wish to specify a threshold for the probability of false alarm. Thus, we instead use as a metric of discriminatory capability the area under the ROC curve[6] (Receiver Operator Characteristic) based on plots of (PD,PF) for various choices of a level of phosphorylated ZAP70 which sets the threshold for activation. We can think of this as a plot of (PD,PF) for every possible deterministic decision rule a T-cell can make, where a decision rule is uniquely defined by a threshold of activated ZAP70 that must be reached before an immune response is mounted.

Appendix II.

One of the reasons the ZAP70-opening model (referred to in the text as the new model) has a discriminatory advantage over the old model is that the association of Lck with ZAP70 through its SH2 domain provides an extra degree of stabilization to the entire complex, and the effective rate at which any species leaves the complex is slower **(Figure S5)**. Presumably a greater degree of complex-stabilization will lead to greater discriminatory advantages.

To calculate the effective rate at which various species leave the complex between co-receptor, TCR and pMHC, we built a discrete space continuous time Markov chain for a single complex, calculated the mean first passage time (mfpt) for various species to exit the complex, and calculated the effective rate by 1/mfpt. The complex has a total of three bonding interactions, namely that between the coreceptor and pMHC (interaction 1), between the TCR and the pMHC (interaction 2) and between Lck attached to the coreceptor and Zap70 on the ITAM (interaction 3).

The states of the Markov chain represent the various interactions that are still intact **(Figure S6).** In the figure S6, blue states are states for which 2 bonds are still in tact, and red states are those for which only one bond is in tact (i.e. one species has effectively left the complex). For instance, a transition from (1,2,3) to (1,2) indicates that bond 3 has broken but 1,2 are still in tact. This indicates that the pMHC-TCR bond is still in tact and CD8 is still engaged with the pMHC, but that Lck is no longer associated with ZAP70 through its SH2 domain. At this point, either the Lck-ZAP70 bond will reform (and return to state 1,2,3), or bond 1 or 2 can break. If for example bond 1 breaks next, the Markov chain exits to state 2, meaning that CD8 has effectively left the complex.

As an example of how we can use this Markov chain to calculate the effective rate at which any species will leave the complex, we can calculate the mean first

passage time for the Markov chain beginning at state (1,2,3) to exit to any of the red states. Note that in principle this can happen an infinite number of ways, as the Markov process can bounce between any of the blue states an infinite number of times. To solve for the mean first passage times summing over all possible paths, we solve a standard set of recursive equations. Let τ_i indicate the mean first passage time for a Markov process starting in state i to exit to a red state, and λ_i be the total rate of leaving state i (meaning that the expected time to leave state i is $1/\lambda_i$). As an example:

$$\lambda_{(1,2,3)} = k_{off}^{CD8} + k_{off}^{pMHC} + k_{off}^{Lck}$$

The recursive equations are obtained from:

$$\tau_i = \frac{1}{\lambda_i} + \sum_j P_{i\to j}\tau_j$$

The $P_{i\to j}$ is the probability to transition from state i to j, and is obtained by dividing the rate of transitioning from i->j over the total rate of leaving state i. We do the calculation assuming that on rates within the complex are the same as those in the well mixed state. For instance, the rate from state (2,3) to (1,2,3) is chosen to be the on rate between CD8 and a pMHC that would be recognized if CD8 and pMHC were not bound together in a complex. Undoubtedly the on rate within the complex will be much higher, and therefore the effective off rates from this calculation are likely to be conservative, in that the actual effective stabilization of the complex will be much greater than that predicted from this calculation.

To calculate the effective rate at which a species, for instance CD8, leaves the complex, we weight the rate calculated above by the probability that CD8 leaves first. This is calculated in a similar manner as before by a recursive set of equations. We call P_i the probability that the Markov chain, beginning in state i, first reaches the red state corresponding to (2) before reaching any of the other red states. The recursive equations are

$$P_i = \sum_j P_{i\to j}P_j$$

The above equations are solved subject to appropriate boundary conditions: $P_2 = 1, P_1 = P_3 = 0$. Combining the results of the previous calculations, the effective rate at which CD8 leaves the complex is given by $\tau_{(1,2,3)}P_{(1,2,3)}$. What we find is that this calculation predicts the complex will be enhanced considerably due. The effective off rate of CD8 is reduced by a factor of roughly 4-10 for the range of peptide off rates considered here. We thus make an estimate that CD8 off rates become .25 the values considered in Appendix IV when the full complex is in tact.

Appendix III.

To account for the fact that TCR is able to become sufficiently phosphorylated for a steady state proportion of ZAP70 binding, but ZAP70 is not able to become phosphorylated by Lck, it was suggested that perhaps kinetics of ZAP70 are fast. Perhaps ZAP70 unbinds with a faster rate than Lck phosphorylation, and therefore Lck diffuses from the ITAM before it is able to phosphorylate ZAP70. (We note that the corresponding on-rate of ZAP70 would also have to be high enough for a steady state fraction of ZAP70 to be appreciably bound). An immediate

implication of this model is that, since ZAP70 certainly does become phosphorylated and activated upon TCR engagement with foreign peptide-MHC, the kinetics of ZAP70 must be different in the presence of foreign antigen so that ZAP70 is then able to remain bound long enough for phosphorylation by active Lck.

One way to account for this is to postulate two sets of ZAP70 off rates, one off-rate for ZAP70 from a TCR that is not ligated with pMHC, and another slower off-rate for ZAP70 from TCR that is ligated with pMHC. This will result in ZAP70 unbinding faster on average from the TCR when only self peptide is around (since the TCR is more likely to be un-ligated due to the low binding affinity of self peptide) and ZAP70 unbinding slower on average from the TCR in presence of foreign antigen. We note that there is not significant experimental evidence that such a difference in kinetic off rates exists, nor is there a molecular mechanism by which this can be explained. However, we aim here to explore whether or not such a model could recapitulate the experimental results and also outperform the other two models, and if so, can it do so in a parameter regime that is biologically realistic.

To simulate this model, we hold constant the ZAP70 kinetics when the TCR is ligated (and choose these kinetics to coincide with values used for the main text calculations—see **APPENDIX IV**), but vary the second off rate for ZAP70 from an ITAM that is not ligated. What we find is that the non-monotonicity described for the original model **(FIGURE S3/S4)** disappears. That is, increasing the second off rate for Zap70 from non-ligated ITAM only increases the discrimination **(FIGURE S7/S8)**. The kinetic model outperforms the discriminatory capabilities of the original model. However, we also see that the kinetic model predicts that very little ZAP70 will be bound in the absence of agonist, and much of the ZAP70 bound is still phosphorylated **(FIGURE S9)**. Therefore the kinetic model cannot replicate experimental results. Note that Figure S9 shows results for model of TCR signaling in which Lck must undergo an activation step.

We also find that, for the kinetic model to outperform the original model in terms of signal discrimination requires a 10-20x difference in kinetic off rates for ZAP70 from ligated and non-ligated TCR.

We also compared the energetic costs for both the kinetic and original model at steady state in the presence of only self peptide; this is the relevant value to compare (as opposed to energetic consumption in the presence of non-self, or some weighted combination of the two) since the T-cell will spend the vast majority of its time exposed to endogenous peptide. What we find is that increasing the second off rate monotonically increases the energetic consumption over the range of values considered **(FIGURE S10)**.

An intuitive explanation for this is that driving up the second off rate of ZAP70 leaves the ITAM exposed and requires more frequent phosphorylation and dephosphorylation. The technical explanation is that dephosphorylation of ITAM in the absence of coreceptor occurs at a much faster rate than the reverse process, phosphorylation of ITAM by free Lck not attached to coreceptor. Increasing the second off rate of ZAP70 and leaving ITAM open to more frequent dephosphorylations means that a greater frequency of irreversible reactions must occur at steady state, leading to a greater energetic cost. On the other hand, in the limit of very strong ZAP70 binding affinity, the network will look like a simple birth

death process tracking the number of phosphorylated ZAP70. Such networks are at equilibrium and satisfy detailed balance, meaning they require no energy at steady state. This suggests that the kinetic model outperforms the original model in discrimination, but does so at an energetic cost, with greater discrimination requiring greater energetic expenditure.

Appendix IV.

To simulate the protein interaction networks outlined above, we use the Gillespie algorithm for simulating Markov jump processes[1-3]. The Gillespie algorithm is based on the fact that time intervals between chemical reactions are exponentially distributed with a rate equal to the sum of all reaction events that can occur given the current state of the system (namely, the copy number of all molecules in the simulation volume). The system is updated by drawing from an exponential distribution to determine the time of the next reaction, and a multinomial distribution to determine the particular reaction that occurred.

Stochastic simulations of the membrane APC interface are carried out in a 1 micron by 1 micron simulation box. Simulations are carried out using both well-mixed and diffusive (spatially resolved) settings. For spatially resolved calculations, the box is divided into 10,000 boxes of .01 micron squared area. Parameter values listed in the table below carry 'bare' units of 1/s, but can easily be converted to conventional units of micromolar using a subvolume edge of .01 micron. Values of k_on in the below tables are for well mixed simulations. For spatially resolved calculations, these must be multiplied by 10,000.

The reaction scheme involves TCR+pMHC+CD8 complex formation. Figures in the main text show results in which Lck only phosphorylates at a single rate, and in Figure 6 we show the effect of modulating the amount of coreceptor with active Lck. Calculations have also been done allowing Lck to be initially inactive, and requiring an additional step of Lck activation (Figures S9-S12). For calculations with an Lck activation step, we allow Lck activation when CD8 is in the complex, and deactivation occurs only off of the complex. Once TCR has been phosphorylated twice, ZAP70 can bind to the ITAM. After binding, ZAP70 undergoes an 'opening' step (with an opening rate that is varied in the main text). Once open, ZAP70 can become activated by Lck.

Note that, though not specified below, there is an implicit multiplication of 2 for phosphorylation events of TCR that are not phosphorylated at all, and for dephosphorylation of doubly phosphorylated ITAM.

Table 1: Reactions and rate constants used

Reaction	**k_{on}**	**k_{off}**	**k_{cat}**	**Reference**
1. TCR+pMHC	.0052	Varied		7
2. TCR+CD8	.1	20		7
3. Lck activation			1	7
4. Phosphorylation by basal/active Lck			.05/6.26	7/8
5. Lck deactivation			.15	7

6. ITAM dephosphorylation			.05	7
7. Zap70	.0075/Varied	1/Varied		
8. Diffusion			50	

Table 2: Initial conditions

Molecule	Concentration	Reference
Ligand (self+nonself)	100	9
TCR	700	9
CD8	266	9
Zap70	300	9

It is found that spatial resolution does not change the qualitative results of the calculations. All figures with the exception of **Figure S11** are shown for well-mixed systems. **Figure S11** (done with spatial resolution) demonstrates how the qualitative trend is unchanged when including diffusion.

Appendix V.

Entropy production/energy consumption is calculated by keeping a running sum of the natural log of the ratio of forward reaction propensity over reverse reaction propensity for each reaction event that occurs. By reaction propensity, we mean the reaction rate constant times combinatorial factors related to protein number. Due to reference 10, it is shown that the slope of the time average of this quantity equals the rate of entropy production for non equilibrium steady states. To calculate entropy production, it is necessary to include reverse reactions for all involved steps. We choose rate constants to be sufficiently small for reactions not listed above that should not occur in the time scales of a single trajectory such that they are not observed (for instance, the phosphorylation rate by free Lck not associated with coreceptor is chosen such that we do not observe any of these reaction events in the course of a single trajectory—phosphorylation by free Lck serves as the reverse reaction to dephosphorylation of a TCR not associated with CD8-Lck/pMHC complex). Varying the reverse rate of these reactions between 10e-5 and 10e-20 shifts only the baseline entropy production but does not change the qualitative trends.

As long as the rate of the composite process of Lck-mediated phosphorylation of Tyrosine 319 on Zap70 and its subsequent binding via its SH2 domain is slower than the rate of dissociation of the TCR from self ligands, the qualitative results discussed in the main text are valid regarding the new model for Zap70 activation. In this range of parameter space, the new model consumes less energy at steady state than the conventional model **(Figure S14)**. This result can be understood intuitively by noting that in the conventional model, if Zap70 is bound to ITAMs upon stimulation by self ligands, a large fraction of bound Zap70 is phosphorylated. This means that, in the conventional model, there are many more phosphorylation and dephosphorylation events, which consume ATP. A more

technical explanation for the entropy production (measure of energy consumption in non-equilibrium systems) is that the new model behaves more like a birth death process in which most Zap70 molecules are bound and the state space counts the number of Zap70 molecules that are phosphorylated. Birth death processes are exactly reversible and thus have no entropy production at steady state.

Additionally, we see in Figure S13 that by doubling the amount of active Lck not associated with co-receptor, the qualitative results in the main text hold.

SI figure captions:

Figure S1: A very simple 4 state markov chain was analyzed to explain non-monotonicity in signal discrimination of the null model (referred to in the text as 'old model') as a function of ZAP-70's unbinding rate. We find that the likelihood ratio of non-self over self (calculated using the steady state probability of occupying state 3) is non-monotonic as a function of ZAP-70 off-rate.

Figure S2: Using the same 4 state markov chain as before, one can reframe the TCR signaling as a binary asymmetric channel, and calculate the channel capacity. We see that the non-monotonicity observed in the likelihood ratio is present in the channel capacity as well.

Figure S3: These are results from 1000 trajectories of length 500s averaged at every second over a burn time of 100s. The off rate for self is 10s^(-1) and for non-self is 1s^(-1), and 2% of the 100 total peptide is non-self. The AUROC is shown for various values of kinetic parameters of ZAP-70 binding/unbinding. We see that the null model demonstrates a non-monotonicity in the kinetic parameters of ZAP-70. For a fixed on rate, the signal discrimination as a function of off rate is non monotonic, and exhibits an optimal value.

Figure S4: We plot the same results as Figure S3 on a 3D axis to demonstrate the non monotonicity.

Figure S5: To calculate the effective off rate of coreceptor from the TCR complex, we build a markov chain whose state space represents the interactions that are currently in tact in the complex. Here, we number the interactions as 1-3, between CD8, TCR and MHC.

Figure S6: Using the numbering convention from Figure S5, we build a markov chain whose state space represents the interactions that are currently in tact. The system begins in state (1,2,3) with all interactions in tact. A transition from (1,2,3) to (1,3) indicates that the pMHC-TCR bond has broken, but coreceptor/TCR interaction and Lck/ZAP-70 interactions are still in tact.

Figure S7: Here we use calculations similar to figure S3/S4 but for a 'kinetic model' in which ZAP-70 has two sets of kinetic off rates. The off rate when TCR is ligated is

fixed at .5s^(-1), and the off-rate/on-rate for ZAP-70 from unligated TCR is varied. This demonstrates that for a fixed off rate when TCR is ligated, increasing the off rate when TCR is not ligated will only serve to increase the discriminatory capabilities.

Figure S8: We plot the same results as Figure S7 on a 3D axis.

Figure S9: In panel 1, we see the percent of total ZAP-70 in the simulation that is bound to TCR in steady state in the absence of agonist peptide. The null model and ZAP-opening model have high percent of ZAP-70 bound, whereas the kinetic model has significantly lower bound ZAP-70. In panel 2, we see the percent of bound ZAP-70 that is activated. What we find is that the ZAP-opening model has a much lower fraction of ZAP-70 bound to be activated; this implies that the ZAP-opening model does the best job at recapitulating experimental results that allow for ZAP-70 to be both bound and not phosphorylated at steady state. In panel 3 we compare the percent increase in activated ZAP-70 upon introduction of 2% agonist peptide to the system. We find that both the kinetic model and the ZAP-opening model outperform the null model. However, a very large difference in the two kinetic parameters is required for the kinetic model to outperform the ZAP-opening model. Shown are results from simulations in which the kinetic off rate of ZAP-70 from a non-ligated TCR is 20x larger than the off rate from a ligated TCR.

Figure S10: For the kinetic model, we see that increasing the off rate of ZAP-70 from an unligated TCR increases the energetic costs (entropy production) of the system at steady state.

Figure S11: The average from 150 trajectories each (with self and self+non-self) for spatially resolved calculations to compare the null and Zap opening model. The results are qualitatively unchanged from well-mixed simulations. Spatial resolution serves only to enhance the trends seen before as there is the pronounced effect due to pMHC rebinding.

Figure S12: These figures are to be compared with Figure 4 a-c in the main text. We see that including an explicit step of Lck activation does not change the qualitative results.

Figure S13: After doubling the amount of active Lck not associated with co-receptor, we see the results are qualitatively unchanged.

Figure S14: Energy consumption at steady state in the absence of agonist peptide is minimized in the parameter regime for which $k_{off\ self} > k_{open}$.

1. **Gillespie D.T.** 1976. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. J Comput Phys **22**: 403-434
2. **Gillespie D.T.** 1977. Exact stochastic simulation of coupled chemical reactions. J. Phys. Chem. **25:** 2340-2361

3. **Gillespie D.T.** 1992. A rigorous derivation of the chemical master equation. Physica A **188:** 404-425
4. **Berger J.** 1985. Statistical Decision Theory and Bayesian Analysis (Springer, New York).
5. **Lehmann E.** 1959. Testing Statistical Hypotheses. (John Wiley & Sons, New York).
6. **Fawcett T.** 2006. An introduction to ROC analysis. Pattern Recognition Letters. **27:** 861-874.
7. **Hoerter J, Brzostek J, Artyomov M, Abel S, Casas J, Rybakin V, Ampudia J, Lotz C, Connolly J, Chakraborty A, Gould K, Gascoigne N.** 2013. Coreceptor affinity for MHC defines peptide specificity requirements for TCR interaction with coagonist peptide-MHC. J. Exp. Med. **210:** 1807-1821.
8. **Hui E, Vale Ronald.** 2014. In vitro membrane reconstitution of the T-cell receptor proximal signaling network. Nat. Struc. and mol. Biol. **21:** 133-142.
9. **Stepanek O, Prabhakar A, Osswald C, King C, Bulek A, Naeher D, Beaufils-Hugot M, Abanto M, Galati V, Hausmann B, Lang R, Cole D, Huseby E, Sewell A, Chakraborty AK, Palmer E.** 2014. Coreceptor Scanning by the T Cell Receptor Provides a Mechanism for T Cell Tolerance. Cell. **159:** 333-345.
10. **Lebowitz J, Spohn H.** 1999. A Gallavotti-Cohen-Type Symmetry in the Large Deviation Functional for Stochastic Dynamics. J. Stat. Phys. **95:** 333-365

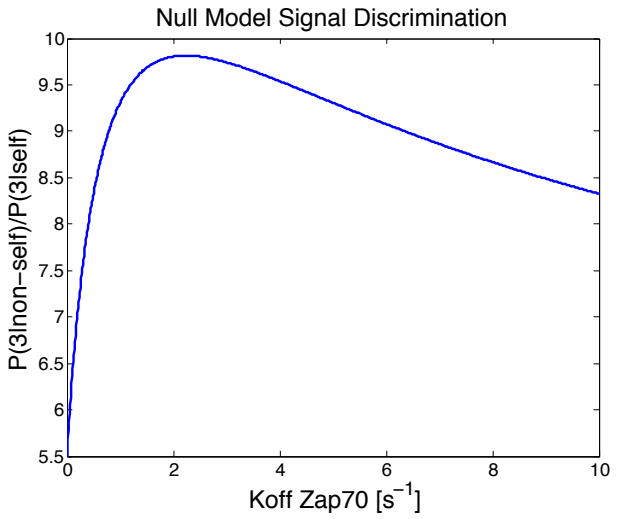

Figure S1

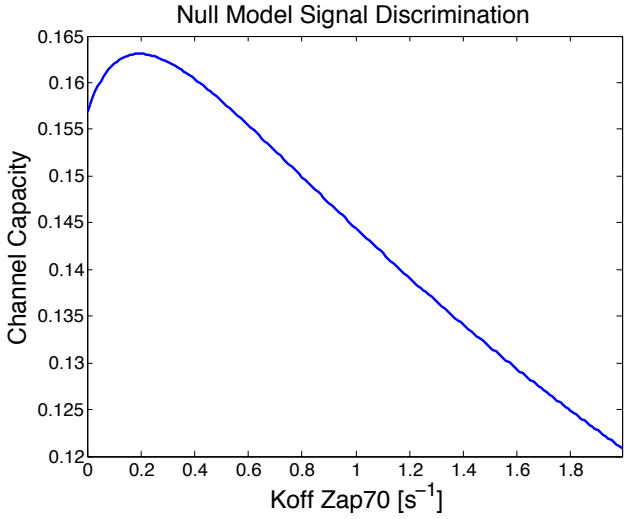

Figure S2

Figure S1: A very simple 4 state markov chain was analyzed to explain non-monotonicity in signal discrimination of the null model (referred to in the text as 'old model') as a function of ZAP-70's unbinding rate. We find that the likelihood ratio of non-self over self (calculated using the steady state probability of occupying state 3) is non-monotonic as a function of ZAP-70 off-rate.

Figure S2: Using the same 4 state markov chain as before, one can reframe the TCR signaling as a binary asymmetric channel, and calculate the channel capacity. We see that the non-monotonicity observed in the likelihood ratio is present in the channel capacity as well.

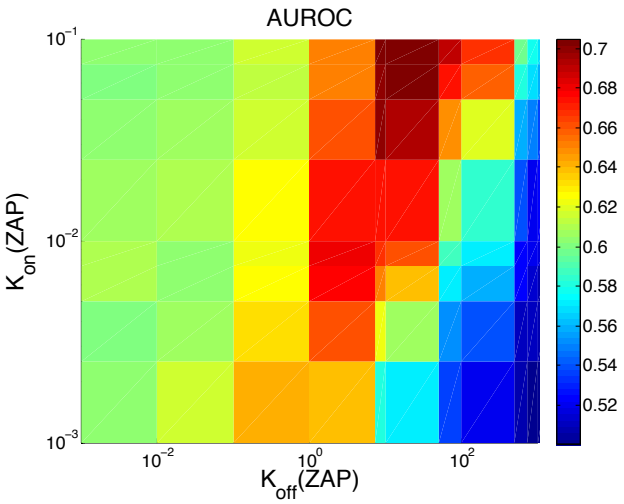

Figure S3

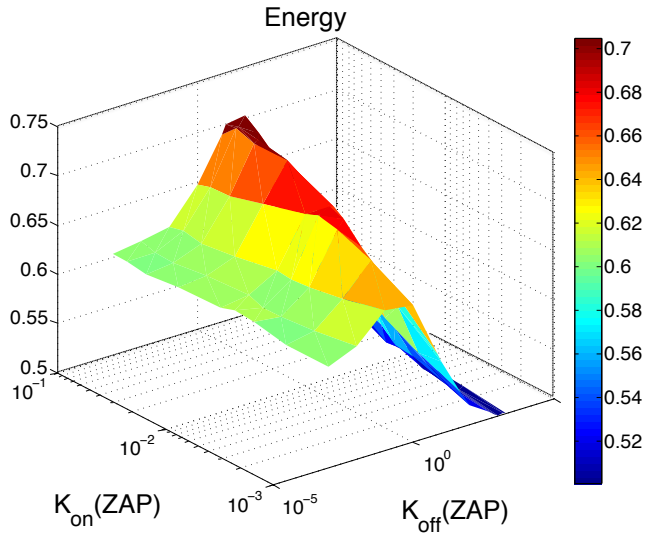

Figure S4

Figure S3: These are results from 1000 trajectories of length 500s averaged at every second over a burn time of 100s. The off rate for self is 10s^(-1) and for non-self is 1s^(-1), and 2% of the 100 total peptide is non-self. The AUROC is shown for various values of kinetic parameters of ZAP-70 binding/unbinding. We see that the null model demonstrates a non-monotonicity in the kinetic parameters of ZAP-70. For a fixed on rate, the signal discrimination as a function of off rate is non monotonic, and exhibits an optimal value.

Figure S4: We plot the same results as Figure S3 on a 3D axes to demonstrate the non monotonicity.

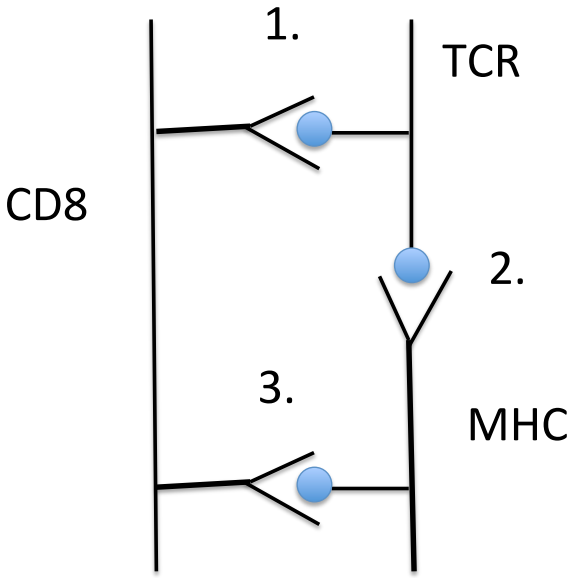

Figure S5

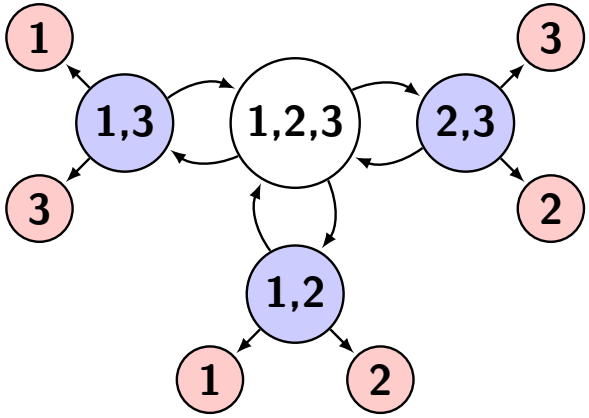

Figure S6

Figure S5: To calculate the effective off rate of coreceptor from the TCR complex, we build a markov chain whose state space represents the interactions that are currently in tact in the complex. Here, we number the interactions as 1-3, between CD8, TCR and MHC.

Figure S6: Using the numbering convention from Figure S5, we build a markov chain whose state space represents the interactions that are currently in tact. The system begins in state (1,2,3) with all interactions in tact. A transition from (1,2,3) to (1,3) indicates that the pMHC-TCR bond has broken, but coreceptor/TCR interaction and Lck/ZAP-70 interactions are still in tact.

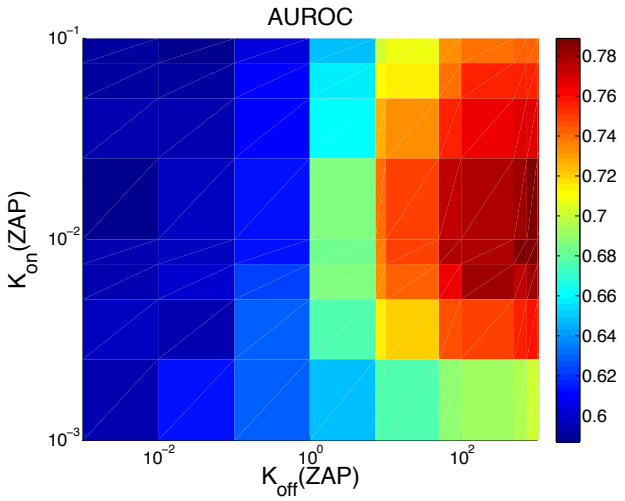

Figure S7

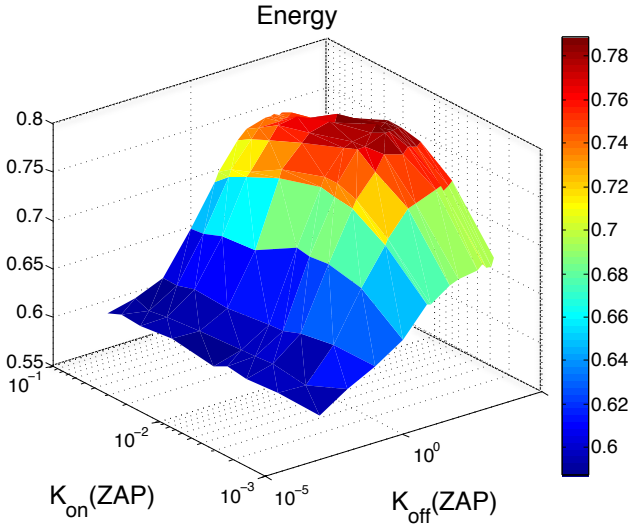

Figure S8

Figure S7: Here we use calculations similar to figure S3/S4 but for a 'kinetic model' in which ZAP-70 has two sets of kinetic off rates. The off rate when TCR is ligated is fixed at .5s^(-1), and the off-rate/on-rate for ZAP-70 from unligated TCR is varied. This demonstrates that for a fixed off rate when TCR is ligated, increasing the off rate when TCR is not ligated will only serve to increase the discriminatory capabilities.

Figure S8: We plot the same results as Figure S7 on a 3D axis.

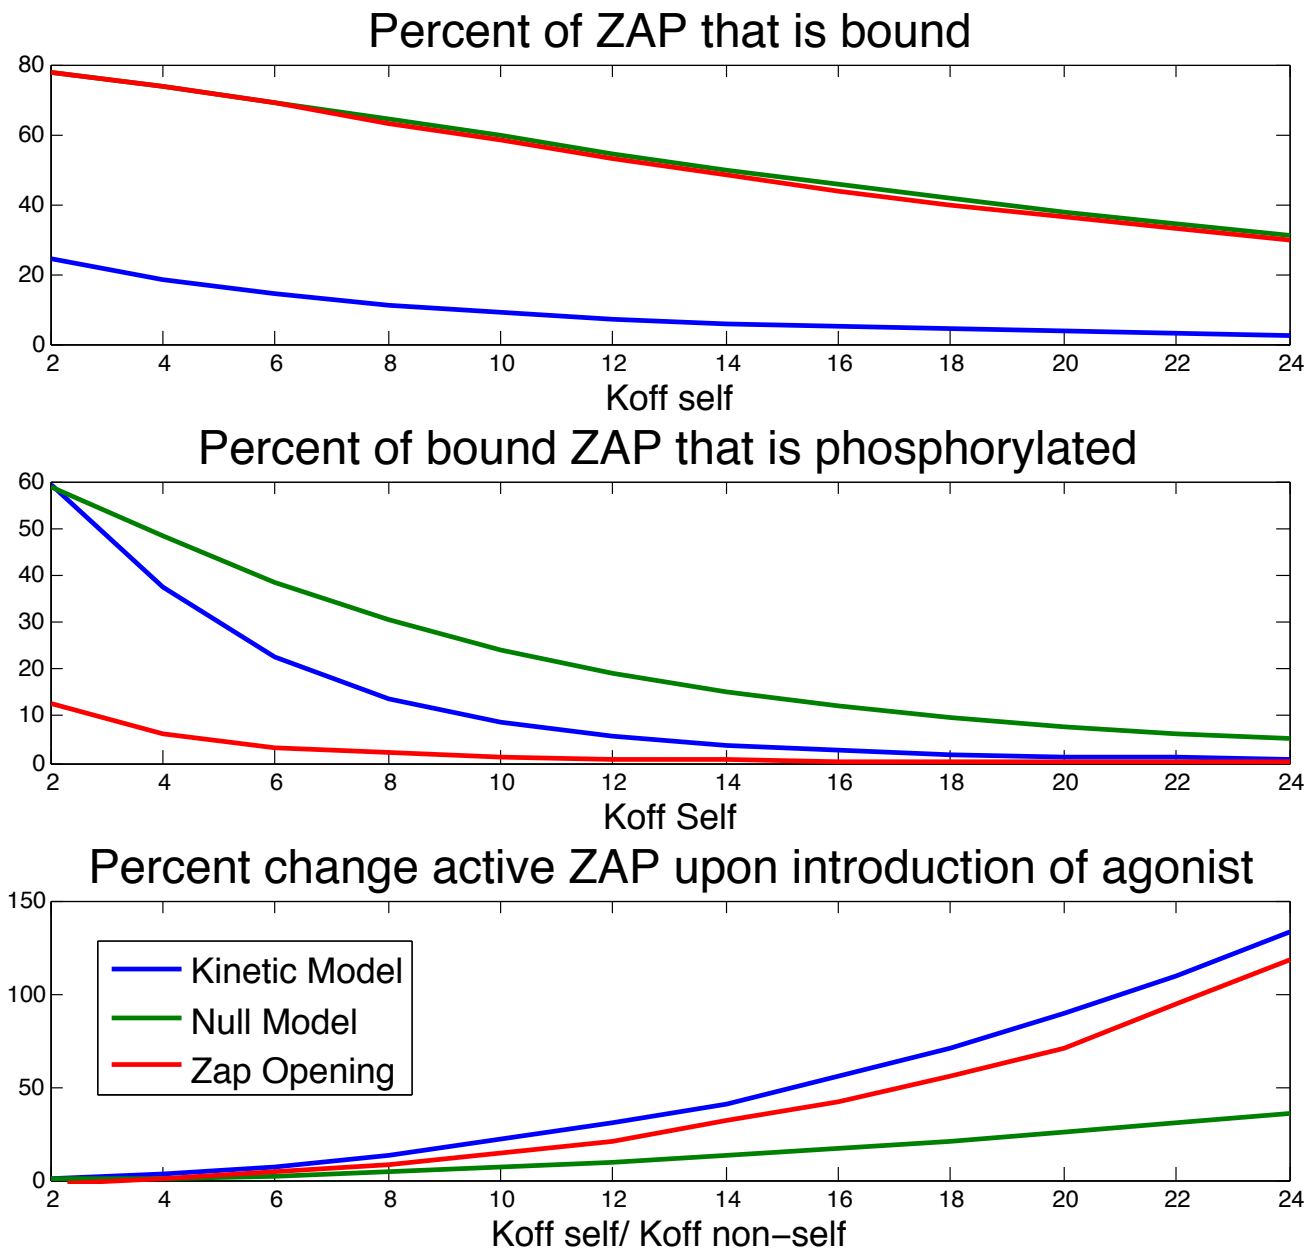
Percent of ZAP that is bound
Koff self
Percent of bound ZAP that is phosphorylated
Koff Self
Percent change active ZAP upon introduction of agonist
Kinetic Model
Null Model
Zap Opening
Koff self/ Koff non−self

Figure S9

Figure S9: Here, we compare various models of TCR signaling, including in the simulation an extra step of Lck activation. In panel 1, we see the percent of total ZAP-70 in the simulation that is bound to TCR in steady state in the absence of agonist peptide. The null model and ZAP-opening model have high percent of ZAP-70 bound, whereas the kinetic model has significantly lower bound ZAP-70. In panel 2, we see the percent of bound ZAP-70 that is activated. What we find is that the ZAP-opening model has a much lower fraction of ZAP-70 bound to be activated; this implies that the ZAP-opening model does the best job at recapitulating experimental results that allow for ZAP-70 to be both bound and not phosphorylated at steady state. In panel 3 we compare the percent increase in activated ZAP-70 upon introduction of 2% agonist peptide to the system. We find that both the kinetic model and the ZAP-opening model outperform the null model. However, a very large difference in the two kinetic parameters is required for the kinetic model to outperform the ZAP-opening model. Shown are results from simulations in which the kinetic off rate of ZAP-70 from a non-ligated TCR is 20x larger than the off rate from a ligated TCR.

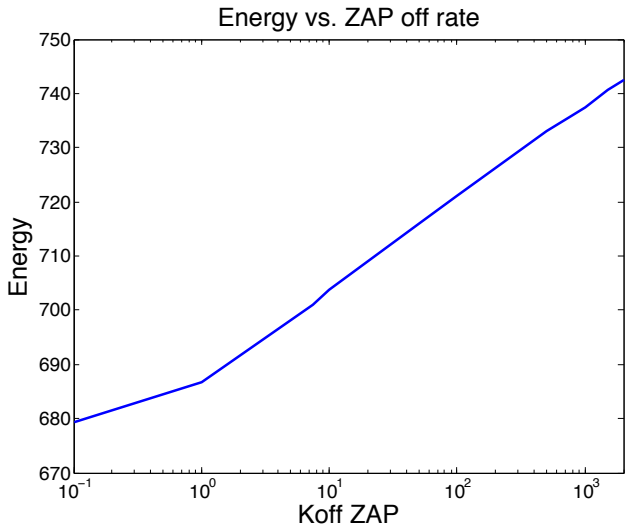

Figure S10

Figure S10: For the kinetic model, we see that increasing the off rate of ZAP-70 from an unligated TCR increases the energetic costs (entropy production) of the system at steady state.

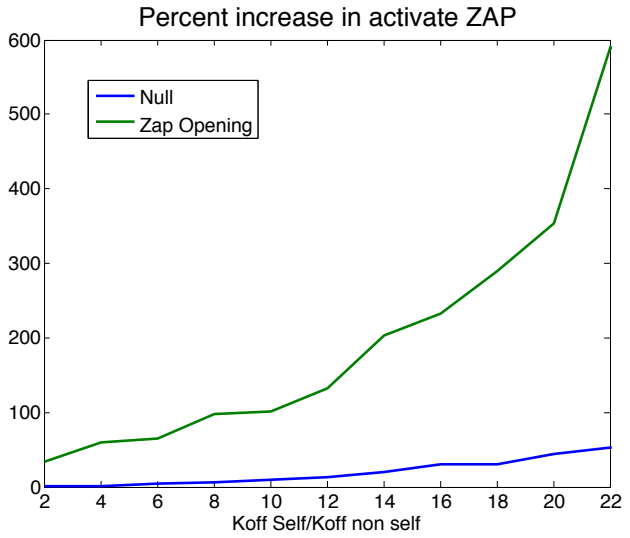

Figure S11

Figure S11: The average from 150 trajectories each (with self and self +non-self) for spatially resolved calculations to compare the null and Zap opening model. The results are qualitatively unchanged from well-mixed simulations. Spatial resolution serves only to enhance the trends seen before as there is the pronounced effect due to pMHC rebinding.

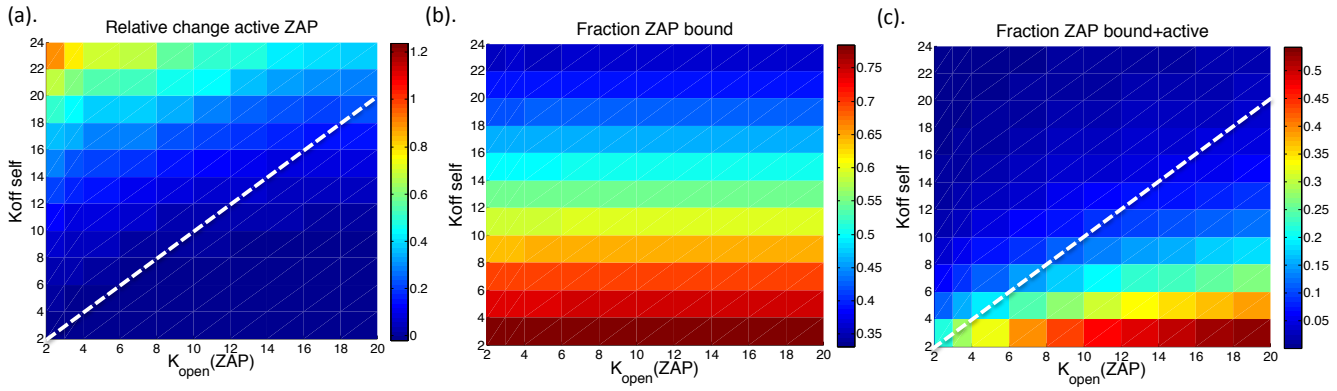

Figure S12: These figures are to be compared with Figure 4 a-c in the main text. We see that including an explicit step of Lck activation does not change the qualitative results.

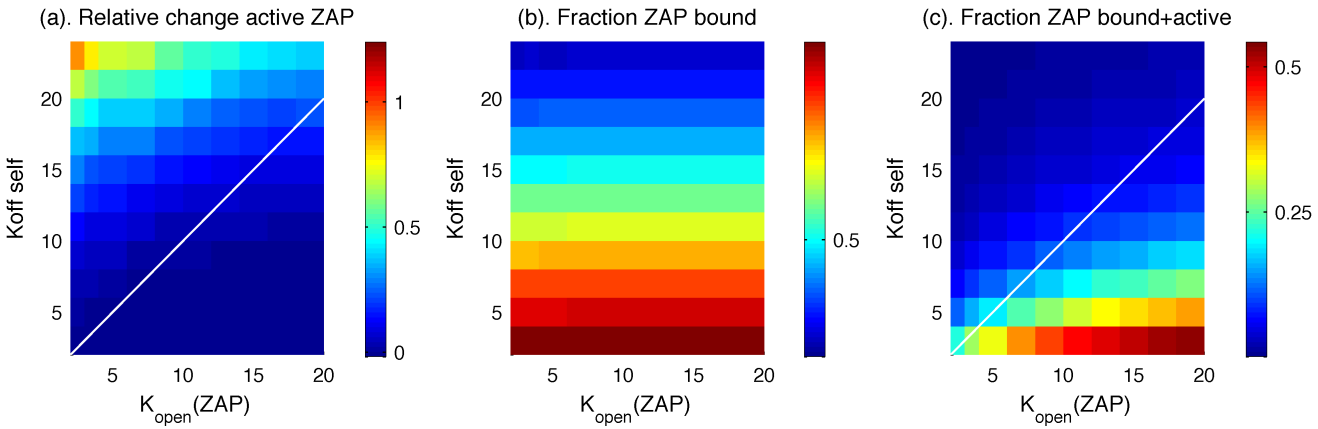

Figure S13: After doubling the amount of active Lck not associated with co-receptor, we see the results are qualitatively unchanged.

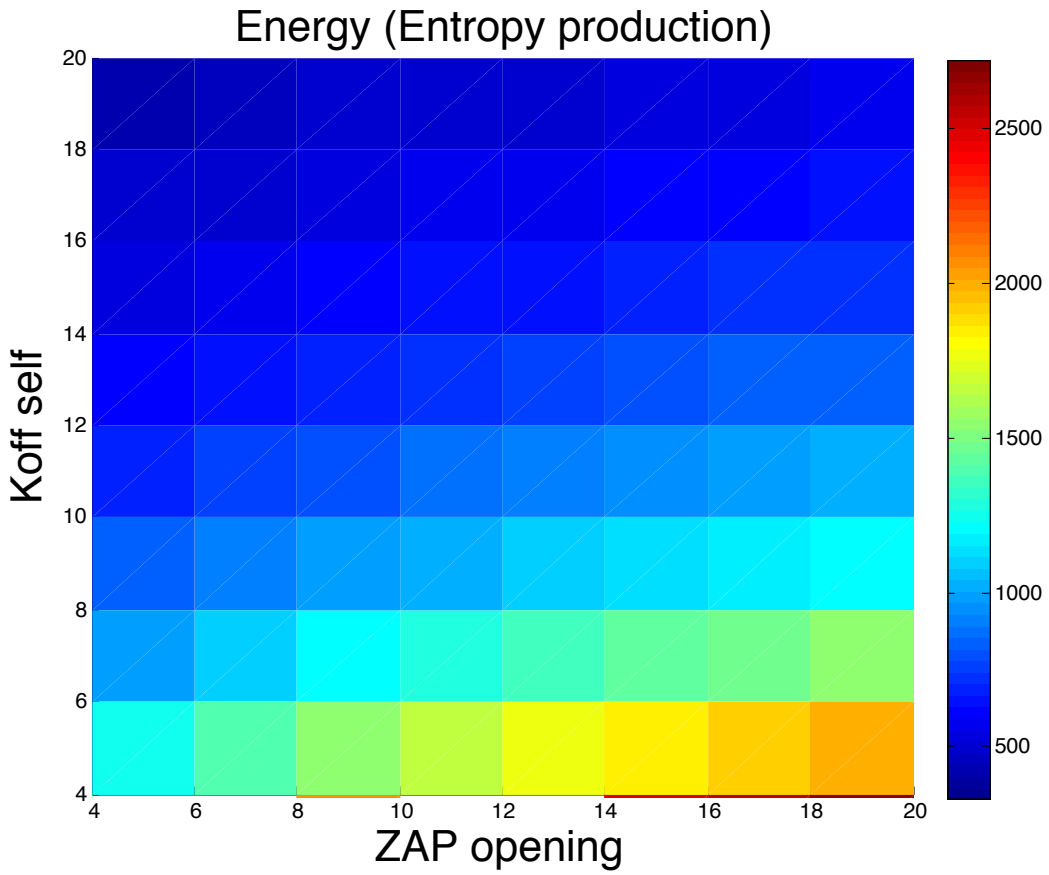

Figure S14: Energy consumption at steady state in the absence of agonist peptide is minimized in the parameter regime for which $k_{off\ self} > k_{open}$.