Supporting Information Text

1. Derivation of the Simple Repression Input-Output Function

In this section, we derive the input-output function for the inducible simple repression motif. This section summarizes the results from Garcia and Phillips 2011 (1) and Razo-Mejia et al. 2018 (2) and we direct the reader to these references for further detail.

We begin by defining the simple repression motif as a regulatory architecture in which binding of a repressor to its cognate binding site occludes binding of an RNA polymerase (RNAP) to the promoter, thereby hindering gene expression (3, 4). The repressor in this work is considered to be an allosteric molecule which fluctuates between an active and inactive state in thermal equilibrium. Binding of an allosteric effector molecule (i.e. an inducer) to a binding site in one domain of the repressor can stabilize the inactive state relative to the active state. The repressor can still bind the DNA in the inactive state but the sequence specificity is reduced and binding to the cognate sequence becomes comparable to nonspecific binding. Such regulatory motifs have been well characterized from a thermodynamic perspective in which the system is considered to be in equilibrium at the time scales relevant to molecular binding events. Under such a model, we can assume that the level of gene expression is proportional to the the probability of RNAP being bound to the promoter and this has frequently been applied in thermodynamic models of transcription (1–10).

The probability of the promoter being occupied by either a polymerase, repressor, or neither is dictated by the Boltzmann distribution,

$$P_{\text{state}} \propto e^{-\varepsilon_{\text{state}}/k_B T},\tag{1}$$

where $\varepsilon_{\text{state}}$ is the energy of the state of interest. This energy is scaled to the thermal energy of the system $k_B T$ where k_B is Boltzmann's constant and T is the temperature in units of K. The goal of this section is to translate generic proportionality in Eq. (1) into the relevant states of our system.

The occupancy states of the promoter and corresponding statistical weights can be seen in Fig. S1 (A). Here we use *P* to denote the number of RNAP per cell, R_A as the number of active repressors, and R_I as the number of inactive repressors. We assume there is a single specific binding site on the genome and N_{NS} nonspecific binding sites. The polymerase, active, and inactive repressor in this work are considered to bind the DNA with different strengths. The energy of binding for each species is given as $\Delta \varepsilon_P$, $\Delta \varepsilon_{RA}$, or $\Delta \varepsilon_{RI}$, which captures the energetic difference between nonspecific and specific binding for that species. Since we consider a single specific binding site for both the polymerase and repressor, we can say that $N_{NS} >> P$ and $N_{NS} >> R_A + R_I$, allowing the multiplicity of arranging *P* polymerases and *R* repressors to be approximately equal to P/N_{NS} and $(R_A + R_I)/N_{NS}$, respectively.



Fig. S1. States and statistical weights for a simple repression motif with an allosteric repressor. (A) Occupancy states of the promoter in which the statistical weights are relative to the unoccupied state. *P*, R_A , and R_I correspond to the average number of RNAP, active repressors, and inactive repressors per cell, respectively. The relative DNA binding energies are given as $\Delta \varepsilon_{P}$, $\Delta \varepsilon_{RA}$ and $\Delta \varepsilon_{RI}$. (B) Allosteric states of the repressor and the statistical weights relative to the active repressor with no inducer bound. The dissociation constant for the inducer to the active and inactive state of the repressor are denoted as K_A and K_I , respectively, *c* is the inducer concentration, and $\Delta \varepsilon_{AI}$ is the energetic difference between the active and inactive states of the repressor.

With these states and statistical weights enumerated, we can now define the probability of a polymerase being bound to the promoter as

$$P_{\text{bound}} = \frac{\frac{P}{N_{NS}}e^{-\beta\Delta\varepsilon_{P}}}{1 + \frac{P}{N_{NS}}e^{-\beta\Delta\varepsilon_{P}} + \frac{R_{A}}{N_{AS}}e^{-\beta\Delta\varepsilon_{RA}} + \frac{R_{I}}{N_{NS}}e^{-\beta\Delta\varepsilon_{RI}}},$$
[2]

where we have defined β as $\frac{1}{k_{\rm P}T}$.

It is experimentally difficult to measure P_{bound} directly as identifying the direct proportionality to gene expression is not straightforward. However, we can easily measure the fold-change in gene expression, defined as the probability of a polymerase

2 of 37 Griffin Chure, Manuel Razo-Mejia, Nathan M. Belliveau, Tal Einav, Stephanie L. Barnes, Mitchell Lewis, Rob Phillips

being bound to the promoter in the presence of repressor relative to constitutive expression,

$$\text{fold-change} = \frac{P_{\text{bound}}(R>0)}{P_{\text{bound}}(R=0)} = \frac{1 + \frac{P}{N_{NS}}e^{-\beta\Delta\varepsilon_P}}{1 + \frac{P}{N_{NS}}e^{-\beta\Delta\varepsilon_P} + \frac{R_A}{N_{NS}}e^{-\beta\Delta\varepsilon_{RA}} + \frac{R_I}{N_{NS}}e^{-\beta\Delta\varepsilon_{RI}}}.$$
[3]

This can be simplified by making two well justified approximations. We can assume that binding of the inactive repressor to the specific binding site is approximately equal to nonspecific binding, $\frac{R_I}{N_{NS}}e^{-\beta\Delta\epsilon_{RI}} << 1 + \frac{R_A}{N_{NS}}e^{-\beta\Delta\epsilon_{RA}}$. Secondly, we can state that binding of RNAP to the promoter is weak, $\frac{P}{N_{NS}}e^{-\beta\Delta\epsilon_{P}} << 1$. Assuming $P \approx 10^3$ (11), $N_{NS} \approx 4.6 \times 10^6$ (the length of the *E. coli* genome in base pairs) and $\Delta\epsilon_P \approx -2$ to $-5k_BT$ (12), the probability of this state comes to $\approx 1\%$ and can be neglected. Using these approximations, we can state that the fold-change in gene expression has the form

fold-change =
$$\left(1 + \frac{R_A}{N_{NS}}e^{-\beta\Delta\varepsilon_{RA}}\right)^{-1}$$
. [4]

In order to make falsifiable predictions, we must have a precise knowledge of the number of active repressors in the cell R_A . While determining this quantity is fraught with experimental difficulties, it is relatively easy to determine the total number of repressors per cell R through quantitative western blotting (1), fluorescence based methods (9), or proteomic studies (13). We can compute the number of active repressors at a given inducer concentration c by multiplying the total number of repressors by the probability of a repressor being active at that inducer concentration,

$$R_A(c) = p_{\rm act}(c)R.$$
[5]

Similarly to computing P_{bound} , we can compute the possible states and statistical weights of the repressor activity, shown in Fig. S1 (B). Following the model of Monod, Wyman, and Changeux (14), we have defined all statistical weights relative to the active repressor with no bound inducer molecules. We have defined the dissociation constant of the inducer to the active and inactive repressor as K_A and K_I , respectively, and have assigned an energetic penalty $e^{-\beta\Delta\varepsilon_{AI}}$ to all inactive states of the repressor. The energetic term $\Delta\varepsilon_{AI}$ represents the relative energy difference between the active and inactive states, $\Delta\varepsilon_{AI} = \varepsilon_I - \varepsilon_A$. For the *lac* repressor used in this work, the value of $\Delta\varepsilon_{AI}$ has been inferred to be $4.5 k_B T$, indicating that the active state is energetically preferred and with no inducer, approximately 99% of the repressors are in the active state.

Using these states and weights, we can compute $P_{act}(c)$ as

$$p_{\rm act}(c) = \frac{\left(1 + \frac{c}{K_A}\right)^2}{\left(1 + \frac{c}{K_A}\right)^2 + e^{-\beta\Delta\varepsilon_{AI}} \left(1 + \frac{c}{K_I}\right)^2}.$$
[6]

Using Eq. (4) - Eq. (6), we can then state that at a given inducer concentration *c*, the fold-change in gene expression can be defined as

fold-change =
$$\left(1 + \frac{\left(1 + \frac{c}{K_A}\right)^2}{\left(1 + \frac{c}{K_A}\right)^2 + e^{-\beta\Delta\varepsilon_{AI}} \left(1 + \frac{c}{K_I}\right)^2} \frac{R}{N_{NS}} e^{-\beta\Delta\varepsilon_{RA}}\right)^{-1},$$
[7]

which is the result stated in Eq. 1 of the main text. We emphasize that equilibrium models as derived here have frequently been used to characterize the simple repression motif (1-5, 9) in addition to non-equilibrium approaches which have the same functional form as Eq. (7) (15).

2. Bayesian Parameter Estimation For DNA Binding Mutants

In this section, we outline the statistical model used in this work to estimate the DNA binding energy for a given mutation in the DNA binding domain. We begin with a derivation of our statistical model using Bayes' theorem and then perform a series of principled steps to validate our choices of priors, ensure computational feasibility, and assess the validity of the model given the collected data. This work follows the analysis pipeline outlined by Michael Betancourt in his case-study entitled "Towards A Principled Bayesian Workflow."

The second subsection "Building a Generative Statistical Model" lays out the statistical model used in this work to estimate the DNA binding energy and the error term σ . The subsequent subsections – "Prior Predictive Checks", "Simulation Based Calibration", and "Posterior Predictive Checks" – define and summarize a series of tests that ensure that the parameters of the statistical model can be identified and are computationally tractable. To understand how we defined our statistical model, only the second subsection is needed.

Calculation of the Fold-Change in Gene Expression. We appreciate the subtleties of the efficiency of photon detection in the flow cytometer, fluorophore maturation and folding, and autofluorescence correction, and we understand the importance in modeling the effects that these processes have on the reported value of the fold-change. However, in order to be consistent with the methods used in the literature, we took a more simplistic approach to calculate the fold-change. Given a set of fluorescence measurements of the constitutive expression control (R = 0), an autofluorescence control (no YFP), and the experimental strain (R > 0), we calculate the fold-change as

$$fold-change = \frac{\langle I_{cell}(R>0) \rangle - \langle I_{autofluorescence} \rangle}{\langle I_{cell}(R=0) \rangle - \langle I_{autofluorescence} \rangle}.$$
[8]

It is important to note here that for a given biological replicate, we consider only a point estimate of the mean fluorescence for each sample and perform a simple subtraction to adjust for background fluorescence. For the analysis going forward, all mentions of measured fold-change are determined by this calculation.

Building a Generative Statistical Model. To identify the minimal parameter set affected by a mutation, we assume that mutations in the DNA binding domain of the repressor alters only the DNA binding energy $\Delta \varepsilon_{RA}$, while the other parameters of the repressor are left unperturbed from their wild-type values. As a first approach, we can assume that all of the other parameters are known without error and can be taken as constants in our physical model. Ultimately, we want to know how probable a particular value of $\Delta \varepsilon_{RA}$ is given a set of experimental measurements *y*. Bayes' theorem computes this distribution, termed the *posterior distribution* as

$$g(\Delta \varepsilon_{RA} | y) = \frac{f(y | \Delta \varepsilon_{RA}) g(\Delta \varepsilon_{RA})}{f(y)},$$
[9]

where we have used g and f to represent probability densities over parameters and data, respectively. The expression $f(y | \Delta \varepsilon_{RA})$ captures the likelihood of observing our data set y given a value for the DNA binding energy under our physical model. All knowledge we have of what the DNA binding energy *could* be, while remaining completely ignorant of the experimental measurements, is defined in $g(\Delta \varepsilon_{RA})$, referred to as the *prior distribution*. Finally, the likelihood that we would observe the data set y while being ignorant of our physical model is defined by the denominator f(y). In this work, this term serves only as a normalization factor and as a result will be treated as a constant. We can therefore say that the posterior distribution of $\Delta \varepsilon_{RA}$ is proportional to the joint distribution between the likelihood and the prior,

$$g(\Delta \varepsilon_{RA} \mid y) \propto f(y \mid \Delta \varepsilon_{RA}) g(\Delta \varepsilon_{RA}).$$
^[10]

We are now tasked with translating this generic notation into a concrete functional form. Our physical model given by Eq. (7) computes the average fold-change in gene expression. Speaking practically, we make several replicate measurements of the fold-change to reduce the effects of random errors. As each replicate is independent of the others, it is reasonable to expect that these measurements will be normally distributed about the theoretical value of the fold-change μ , computed for a given $\Delta \varepsilon_{RA}$. We can write this mathematically for each measurement as

$$f(y \mid \Delta \varepsilon_{RA}) = \frac{1}{(2\pi\sigma^2)^{N/2}} \prod_{i}^{N} \exp\left[\frac{-(y_i - \mu(\Delta \varepsilon_{RA}))^2}{2\sigma^2}\right],$$
[11]

where *N* is the number of measurements in *y* and y_i is the i^{th} experimental fold-change measurement. We can write this likelihood in shorthand as

$$f(y \mid \Delta \varepsilon_{RA}) = \text{Normal}\{\mu(\Delta \varepsilon_{RA}), \sigma\}$$
[12]

which we will use for the remainder of this section.

Using a normal distribution for our likelihood has introduced a new parameter σ which describes the spread of our measurements about the true value. We must therefore include it in our parameter estimation and assign an appropriate prior distribution such that the posterior distribution becomes

$$g(\Delta \varepsilon_{RA}, \sigma | y) \propto f(y | \Delta \varepsilon_{RA}, \sigma) g(\Delta \varepsilon_{RA}) g(\sigma).$$
[13]

We are now tasked with assigning functional forms to the priors $g(\Delta \varepsilon_{RA})$ and $g(\sigma)$. Though one hopes that the result of the inference is not too dependent on the choice of prior, it is important to choose one that is in agreement with our physical and physiological intuition of the system.

We can impose physically reasonable bounds on the possible values of the DNA binding energy $\Delta \varepsilon_{RA}$. We can say that it is unlikely that any given mutation in the DNA binding domain will result in an affinity greater than that of biotin to streptavidin [1 fM $\approx -35 k_BT$, BNID 107139 (16)], one of the strongest known non-covalent bonds. Similarly, it's unlikely that a given mutation will result in a large, positive binding energy, indicating non-specific binding is preferable to specific binding (~ 1 to 10 k_BT). While it is unlikely for the DNA binding energy to exceed these bounds, it's not impossible, meaning we should not impose these limits as hard boundaries. Rather, we can define a weakly informative prior as a normal distribution with a mean and standard deviation as the average of these bounds,

$$g(\Delta \varepsilon_{RA}) \sim \text{Normal}\{-12, 12\}$$
 [14]

whose probability density function in shown in Fig. S2 (A).

By definition, fold-change is restricted to the bounds [0, 1]. Measurement noise and fluctuations in autofluorescence background subtraction means that experimental measurements of fold-change can extend beyond these bounds, though not substantially. By definition, the scale parameter σ must be positive and greater than zero. We also know that for the measurements to be of any use, the error should be less than the available range of fold-change, 1.0. We can choose such a prior as a half normal distribution

$$g(\sigma) = \frac{1}{\phi} \sqrt{\frac{2}{\pi}} \exp\left[-\frac{\sigma^2}{2\phi^2}\right]; \forall \sigma \ge 0,$$
[15]

where ϕ is the standard deviation. By choosing $\phi = 0.1$, it is unlikely that $\sigma \ge 1$ yet not impossible, permitting the occasional measurement significantly outside of the theoretical bounds. The probability density function for this prior is shown in Fig. S2(B).

While these choices for the priors seem reasonable, we can check their appropriateness by using them to simulate a data set and checking that the hypothetical fold-change measurements obey our physical and physiological intuition.

Prior Predictive Checks. If our choice of prior distribution for each parameter is appropriate, we should be able to simulate data sets using these priors that match our expectations. In essence, we would hope that these prior choices would generate some data sets with fold-change measurements above 1 or below zero, but they should be infrequent. If we end up getting primarily negative values for fold-change, for example, then we can surmise that there is something wrong in our definition of the prior distribution. This method, coined a *prior predictive check*, was first put forward by Isidore Good in 1950 (17) and has received newfound attention in computational statistics.

We perform the simulation in the following manner. We first draw a random value for $\Delta \varepsilon_{RA}$ out of its prior distribution stated in Eq. (14) and calculate what the mean fold-change should be given our theory described in Eq. (7). With this in hand, we draw a random value for σ from its prior distribution, specified in Eq. (15). We then generate a simulated dataset by drawing \approx 70 fold-change values across twelve inducer concentrations from the likelihood distribution which we defined in Eq. (12). This roughly matches the number of measurements made for each mutant in this work. We repeat this procedure for 800 draws from the prior distributions, which is enough to observe the occasional extreme fold-change value from the likelihood. As the DNA binding energy is the only parameter of our physical model that we are estimating, we had to choose values for the others. We kept the values of the inducer binding constants K_A and K_I the same as the wild-type repressor (139 μ M and 0.53 μ M, respectively). We chose to use R = 260 repressors per cell as this is the repressor copy number we used in the main text to estimate the DNA binding energies of the three mutants.

The draws from the priors are shown in S2(A) and (B) as black points above the corresponding distribution. To display the results, we computed the percentiles of the simulated data sets at each inducer concentration. These percentiles are shown as red shaded regions in Fig. S2(C). The 5th percentile (dark red band) has the characteristic profile of an induction curve. Given that the prior distribution for $\Delta \varepsilon_{RA}$ is centered at $-12 k_B T$ and we chose R = 260, we expect the generated data sets to cluster about the induction profile defined by these values. More importantly, approximately 95% of the generated data sets fall between fold-change values of -0.1 and 1.1, which is within the realm of possibility given the systematic and biological noise in our experiments. The 99th percentile maximum is approximately 1.3 and the minimum approximately -0.3. While we could tune our choice of prior further to minimize draws this far from the theoretical bounds, we err on the side of caution and accept these values as it is possible that fold-change measurements this high or low can be observed, albeit rarely.

Through these prior predictive checks, we feel confident that these choices of priors are appropriate for the parameters we wish to estimate. We can now move forward and make sure that the statistical model as a whole is valid and computationally tractable.

Sensitivity Analysis and Simulation Based Calibration. Satisfied with our choice of prior distributions, we can proceed to check other properties of the statistical model and root out any pathologies lurking in our model assumptions.

To build trust in our model, we could generate a data set \tilde{y} with a *known* value for σ and $\Delta \varepsilon_{RA}$, estimate the posterior distribution $g(\Delta \varepsilon_{RA}, \sigma | \tilde{y})$, and determine how well we were able to retrieve the true value of the parameters. However, running this once or twice for handpicked values of σ and $\Delta \varepsilon_{RA}$ won't reveal edge-cases in which the inference fails, some of which may exist in our data. Rather than performing this operation once, we can run this process over a variety of data sets where the ground truth parameter value is drawn from the prior distribution (as we did for the prior predictive checks). For an arbitrary parameter θ , the joint distribution between the ground truth value $\tilde{\theta}$, the inferred value θ , and the simulated data set \tilde{y} can be written as

$$\pi(\theta, \tilde{y}, \tilde{\theta}) = g(\theta \,|\, \tilde{y}) f(\tilde{y} \,|\, \tilde{\theta}) g(\tilde{\theta}).$$
[16]



Fig. S2. Prior distributions and prior predictive check for estimation of the DNA binding energy. (A) Prior probability density function for DNA binding energy $\Delta \varepsilon_{RA}$ as \sim Normal(-12, 12). (B) Prior probability density function for the standard deviation in measurement noise σ as \sim HalfNormal(0, 0.1). (C) Percentiles of values drawn from the likelihood distribution given draws from prior distributions given R = 260, $K_A = 139 \times 10^{-6}$ M, $K_I = 0.53 \times 10^{-6}$ M, and $\Delta \varepsilon_{AI} = 4.5 k_B T$, which match the parameters used for the predictions in Razo-Mejia et al. 2018 (2). Black points at top of (A) and (B) represent draws used to generate fold-change measurements from the likelihood distribution. Percentiles in (C) generated from 800 draws from the prior distributions. For each draw from the prior distributions, a data set of 70 measurements over 12 IPTG concentrations (ranging from 0 to 5000 μ M) were generated from the likelihood.

If this process is run for a large number of simulations, Eq. (16) can be marginalized over all data sets \tilde{y} and all ground truth values $\tilde{\theta}$ to yield the original prior distribution,

$$\int d\tilde{\theta} \int d\tilde{y}\pi(\theta, \tilde{y}, \tilde{\theta}) = g(\theta).$$
[17]

This result, described by Talts et al. 2018 (18), holds true for *any* statistical model and is a natural self consistency property of Bayesian inference. Any deviation between the distribution of our inferred values for θ and the original prior distribution $g(\theta)$ indicates that either our statistical model is malformed or the computational method is not behaving as expected. There are a variety of ways we can ensure that this condition is satisfied, which we outline below.

Using the data set generated for the prior predictive checks [shown in Fig. S2(C)], we sampled the posterior distribution and compute $\Delta \varepsilon_{RA}$ and σ for each simulation and checked that they matched the original prior distribution. To perform the inference, we use Markov chain Monte Carlo (MCMC) to sample the posterior distribution. Specifically, we use the Hamiltonian Monte Carlo algorithm implemented in the Stan probabilistic programming language (19). The specific code files can be accessed through the paper website or the associated GitHub repository. The original prior distribution and the distribution of inferred parameter values can be seen in Fig. S3 (A) and (B). For both $\Delta \varepsilon_{RA}$ and σ , we can accurately recover the ground truth distribution (blue) via sampling with MCMC (red). For $\Delta \varepsilon_{RA}$, there appears to be an upper and lower limit past which we are unable to accurately infer the binding energy. This can be seen in both the histogram [Fig. S3(A)] and the empirical cumulative distribution [Fig. S3(B)] as deviations from the ground truth when DNA binding is below $\approx -25k_BT$ or above $\approx -5k_BT$. These limits hinder our ability to comment on exceptionally strong or weak binding affinities. However, as all mutants queried in this work exhibited binding energies between these limits, we surmise that the inferential scheme permits us to draw conclusions about the inferred DNA binding strengths.

Rather than examining the agreement of the data-averaged posterior and the ground truth prior distribution solely by eye, we can compute summary statistics using the mean μ and standard deviation σ of the posterior and prior distributions which permit easier identification of pathologies in the inference. One such quantity is the posterior *z*-score, which is defined as

$$z = \frac{\mu_{\text{posterior}} - \theta}{\sigma_{\text{posterior}}}.$$
[18]

This statistic summarizes how accurately the posterior recovers the ground truth value beyond simply reporting the mean, median, or mode of the posterior distribution. Z-scores around 0 indicate that the posterior is concentrating tightly about the true value of the parameter whereas large values (either positive or negative) indicate that the posterior is concentrating elsewhere. A useful feature of this metric is that the width of the posterior is also considered. It is possible that the posterior could have a mean very close to the ground truth value, but have an incredibly narrow distribution/spread such that it does not overlap with the

6 of 37 Griffin Chure, Manuel Razo-Mejia, Nathan M. Belliveau, Tal Einav, Stephanie L. Barnes, Mitchell Lewis, Rob Phillips



Fig. S3. Comparison of averaged posterior and prior distributions for $\Delta \varepsilon_{RA}$ and σ . (A) Distribution of the average values for the DNA binding energy $\Delta \varepsilon_{RA}$ (red) overlaid with the ground truth distribution (blue). (B) Data averaged posterior (red) for the standard deviation of fold-change measurements overlaid with the ground truth distribution (blue). Top and bottom show the same data with different visualizations.

ground-truth. Only comparing the mean value to the ground truth would suggest that the inference "worked". However with a small standard deviation generates a very large *z*-score, telling us that something has gone awry.

If our inferential model is behaving properly, the width of the posterior distribution should be significantly smaller than the width of the prior, meaning that the posterior is being informed by the data. The level to which the posterior is being informed by the data can be easily calculated given knowledge of both the prior and posterior distribution. This quantity, aptly named the shrinkage *s*, can be computed as

$$s = 1 - \frac{\sigma_{\text{posterior}}^2}{\sigma_{\text{prior}}^2}.$$
[19]

When the shrinkage is close to zero, the variance of the posterior is approximately the same as the variance of the prior, model is not being properly informed by the data. When $s \approx 1$, the variance of the posterior is much smaller than the variance of the prior, indicating that the it is being highly informed by the data. A shrinkage less than 0 indicates that the posterior is wider than the prior distribution, revealing a severe pathology in either the model itself or the implementation.

In Fig. S4, we compute these summary statistics for each parameter. For both $\Delta \varepsilon_{RA}$ and σ , we see clustering of the *z*-score about 0 with the extrema reaching $\approx \pm 3$. This suggests that for the vast majority of our simulated data sets, the posterior distribution concentrated about the ground truth value. We also see that for both parameters, the posterior shrinkage *s* is ≈ 1 , indicating that the posterior is being highly informed by the data. There is a second distribution centered ≈ 0.8 for $\Delta \varepsilon_{RA}$, indicating that for a subset of the data sets, the posterior is only $\approx 80\%$ narrower than the prior distribution. These samples are those that were drawn outside of the limits of ≈ -25 to $-5k_BT$ where the inferential power is limited. Nevertheless, the posterior still significantly shrank, indicating that the data strongly informs the posterior.

The general self-consistency condition given by Eq. (17) provides another route to ensure that the model is computationally tractable. Say that we draw a value for the DNA binding energy from the prior distribution, simulate a data set, and sample the posterior using MCMC. The result of this sampling is a collection of N values of the parameter which may be above, below, or equal to the ground-truth value. From this set of values, we select L of them and rank order them by their value. Talts and colleagues (18) derived a general theorem which states that the number of samples less than the ground truth value of the parameter (termed the rank statistic) is uniformly distributed over the interval [0, L]. As Eq. (17) *must* hold true for any statistical model, deviations from uniformity signal that there is a problem in the implementation of the statistical model. How the distribution deviates is also informative as different types of failures result in different distributions. The nature of these deviations, along with a more formal proof of the uniform distribution of rank statistics can be found in Talts et al. 2018 (18) where it was originally derived.

Given the sampling statistics for each of the simulated data sets, we took 800 of the MCMC samples of the posterior distribution



Fig. S4. Inferential sensitivity for estimation of $\Delta \varepsilon_{RA}$ and σ . The posterior *z*-score for each posterior distribution inferred from a simulated data set is plotted against the shrinkage for (A) the DNA binding energy $\Delta \varepsilon_{RA}$ and (B) the standard deviation of fold-change measurements σ

for each of the 800 simulated data sets and computed the rank statistic. The distributions are shown in Fig. S5 as both histograms and ECDFs for the DNA binding energy and standard deviation. The distribution of rank statistics for both parameters appears to be uniform. The gray band overlaying the histograms (top row) as well as the gray envelopes overlaying the ECDFs (bottom row) represent the 99th percentile expected from a true uniform distribution. The uniformity of this distribution, along with the well-behaved *z*-scores and shrinkage for each parameter, tells us that there are no underlying pathologies in our statistical model and that it is computationally tractable. However, this does not mean that it is correct. Whether this model is valid for the actual observed data is the topic of the next section.

Parameter Estimation and Posterior Predictive Checks. We now turn to applying our vetted statistical model to experimental measurements. While the same statistical model was applied to all three DNA binding mutants, here we only focus on the mutant Q21M for brevity.

Using a single induction profile, we sampled the posterior distribution over both the DNA binding energy $\Delta \varepsilon_{RA}$ and the standard deviation σ using MCMC implemented in the Stan programming language. The output of this process is a set of 4000 samples of both parameters along with the value of their log posterior probabilities, which serves as an approximate measure of the probability of each value. The individual samples are shown in Fig. S6. The joint distribution between $\Delta \varepsilon_{RA}$ and σ is shown in the lower left hand corner, and the marginal distributions for each parameter are shown above and to the right of the joint distribution, respectively. The joint distribution is color coded by the value of the log posterior, with yellow and blue corresponding to high and low probability, respectively. The symmetric shape of the joint distribution is a telling sign that there is no correlation between two parameters. The marginal distributions for each parameter are also relatively narrow, with the DNA binding energy covering a range of $\approx 0.6 k_B T$ and σ spanning ≈ 0.02 . To more precisely quantify the uncertainty, we computed the shortest interval of the marginal distribution for each parameter contains 95% of the probability. The bounds of this interval, coined the Bayesian credible region, can accommodate asymmetry in the marginal distribution since the upper and lower bounds of the estimate are reported. In the main text, we reported the DNA binding energy estimated from these data to be $15.43^{+0.06}_{-0.06} k_B T$, where the first value is the median of the distribution and the super- and subscripts correspond to the upper and lower bounds of the credible region, respectively.

While looking at the shape of the posterior distribution can be illuminating, it is not enough to tell us if the parameter values extracted make sense or accurately describe the data on which they were conditioned. To assess the validity of the statistical model in describing actual data, we again turn to simulation, this time using the posterior distributions for each parameter rather than the prior distributions. The likelihood of our statistical model assumes that across the entire induction profile, the observed fold-change is normally distributed about the theoretical prediction with a standard deviation σ . If this is an accurate depiction of the generative process, we should be able to draw values from the likelihood using the sampled values for $\Delta \varepsilon_{RA}$ and σ that are indistinguishable from the actual experimental measurements. This process is known as a *posterior predictive check* and is a Bayesian method of assessing goodness-of-fit.

For each sample from the posterior, we computed the theoretical mean fold-change given the sampled value for $\Delta \varepsilon_{RA}$ and Eq. (7). With this mean in hand, we used the corresponding sample for σ and drew a data set from the likelihood distribution the same size as the real data set used for the inference. As we did this for every sample of our MCMC output (a total of ≈ 4000), it is more instructive to compute the percentiles of the generated data than to show the entire output. In Fig. S6(B), the percentiles of the generated data sets are shown overlaid with the data used for the inference. We see that all of the data points fall within the 99th percentile of simulated data sets with the 5th percentile tracking the mean of the data at each inducer concentration. As there are no systematic deviations or experimental observations that fall far outside those generated from the statistical model, we can safely say that the statistical model derived here accurately describes the observed data.



Fig. S5. Rank distribution of the posterior samples from simulated data. Top row shows a histogram of the rank distribution with n = 20 bins. Bottom row is the cumulative distribution for the same data. Gray bands correspond to the 99th percentile of expected variation from a uniform distribution. (A) Distribution for the DNA binding energy $\Delta \varepsilon_{RA}$ and (B) for the standard deviation σ .



Fig. S6. Markov Chain Monte Carlo (MCMC) samples and posterior predictive check for DNA binding mutant Q21M. (A) Marginal and joint sampling distributions for DNA binding energy $\Delta \varepsilon_{RA}$ and σ . Each point in the joint distribution is a single sample. Marginal distributions for each parameter are shown adjacent to joint distribution. Color in the joint distribution corresponds to the value of the log posterior with the progression of blue to yellow corresponding to increasing probability. (B) The posterior predictive check of model. The measurements of the fold-change in gene expression are shown as black open-faced circles. The percentiles are shown as colored bands and indicate the fraction of simulated data drawn from the likelihood that fall within the shaded region.

3. Inferring the Free Energy From Fold-Change Measurements

In this section, we describe the statistical model to infer the free energy *F* from a set of fold-change measurements. We follow the same principled workflow as described previously for the DNA binding estimation, including declaration of the generative model, prior predictive checks, simulation based calibration, and posterior predictive checks. Finally, we determine an empirical limit in our ability to infer the free energy and define a heuristic which can be used to identify measurements that are likely inaccurate. To understand the statistical model and the empirical limits of detection, only the subsections *Building A Generative Model* and *Sensitivity Limits and Systematic Errors in Inference* are necessary.

Building A Generative Model. In the main text, we showed that the fold-change equation defined in Eq. (7) can be rewritten in the form of a Fermi function,

fold-change =
$$\frac{1}{1 + e^{-F/k_BT}}$$
, [20]

where F corresponds to the free energy difference between the repressor bound and unbound states of the promoter. While the theory prescribes a way for us to calculate the free energy based on our knowledge of the biophysical parameters, we can directly calculate the free energy of a measurement of fold-change by simply rearranging Eq. (20) as

$$F = -k_B T \log\left(\frac{1}{\text{fold-change}} - 1\right).$$
[21]

With perfect measurement of the fold-change in gene expression (assuming no experimental or measurement noise), the free energy can be directly calculated. However, actual measurements of the fold-change in gene expression can extend beyond the theoretical bounds of 0 and 1, for which the free energy is mathematically undefined.

As the fold-change measurements between biological replicates are independent, it is reasonable to assume that they are normally distributed about a mean value μ with a standard deviation σ . While the mean value is restricted to the bounds of [0, 1], fold-change measurements outside of these bounds are still possible given that they are distributed about the mean with a scale of σ . Thus, if we have knowledge of the mean fold-change in gene expression about which the observed fold-change is distributed, we can calculate the mean free energy as

$$F = -k_B T \log\left(\frac{1}{\mu} - 1\right).$$
[22]

For a given set of fold-change measurements y, we wish to infer the posterior probability distribution for μ and σ , given by Bayes' theorem as

$$g(\mu, \sigma | y) \propto f(y | \mu, \sigma)g(\mu)g(\sigma),$$
[23]

where we have dropped the normalization constant f(y) and assigned a proportionality between the posterior and joint probability distribution. Given that the measurements are independent, we define the likelihood $f(y | \mu, \sigma)$ as a normal distribution,

$$f(y \mid \mu \sigma) \sim \text{Normal}\{\mu, \sigma\}.$$
 [24]

While the mean μ is restricted to the interval [0, 1], there is no reason *a priori* to think that it is more likely to be closer to either bound. To remain uninformative and be as permissive as possible, we define a prior distribution for μ as a Uniform distribution between 0 and 1,

$$g(\mu) = \begin{cases} \frac{1}{\mu_{\max} - \mu_{\min}} & \mu_{\min} < \mu < \mu_{\max} \\ 0 & \text{otherwise} \end{cases}.$$
 [25]

Here, $\mu_{\min} = 0$ and $\mu_{\max} = 1$, reducing $g(\mu)$ to 1. For σ , we can again assume a half-normal distribution with a standard deviation of 0.1 as was used for estimating the DNA binding energy [Eq. (15)],

$$g(\sigma) = \text{HalfNormal}\{0, 0.1\}.$$
[26]

With a full generative model defined, we can now use prior predictive checks to ensure that our choices of prior are appropriate for the inference.

Prior Predictive Checks. To check the validity of the chosen priors, we pulled 1000 combinations of μ and σ from their respective distributions [Fig. S7(A)] and subsequently drew a set of 10 fold-change values (a number comparable to the number of biological replicates used in this work) from a normal distribution defined by μ and σ . To visualize the range of values generated from these checks, we computed the percentiles of the empirical cumulative distributions of the fold-change values, as can be seen in Fig. S7(C). Approximately 95% of the the generated fold-change measurements were between the theoretical bounds of [0, 1] whereas 5% of the data sets fell outside with the maximum and minimum values extending to \approx 1.2 and -0.2, respectively. Given our familiarity with these experimental strains and the detection sensitivity of the flow cytometer, these excursions beyond the theoretical bounds agree with our intuition. Satisfied with our choice of prior distributions, we can proceed to check the sensitivity and computational tractability of our model through simulation based calibration.

Griffin Chure, Manuel Razo-Mejia, Nathan M. Belliveau, Tal Einav, Stephanie L. Barnes, Mitchell Lewis, Rob Phillips 11 of 37



Fig. S7. Prior predictive checks for inference of the mean fold-change. (A) The prior distributions for μ (left) and σ (right). The vertical axis is proportional to the probability of the value. Black points above distributions correspond to the values used to perform the prior predictive checks. (B) Percentiles of the data generated for each draw from the prior distributions shown as a cumulative distribution. Percentiles were calculated for 1000 generated data sets, each with 10 fold-change measurements drawn from the likelihood given the drawn values of μ and σ .

Simulation Based Calibration. To ensure that the parameters can be estimated with confidence, we sampled the posterior distribution of μ and σ for each data set generated from the prior predictive checks. For each inference, we computed the *z*-score and shrinkage for each parameter, shown in Fig. S8(A). For both parameters, the *z*-scores are approximately centered about zero, indicating that the posteriors concentrate about the ground truth value of the parameter. The *z*-scores for σ [black points in Fig. S10(A)] appear to be slightly off centered with more negative values than positive. This suggests that σ is more likely to be slightly overestimated in some cases. The shrinkage parameter for μ (red points) is very tightly distributed about 1.0, indicating that the prior is being strongly informed by the data. The shrinkage is more broadly distributed for for σ with a minimum value of ≈ 0.5 . However, the median shrinkage for σ is ≈ 0.9 , indicating that half of the inferences shrank the prior distribution by at least 90%. While we could revisit the model to try and improve the shrinkage values, we are more concerned with μ which shows high shrinkage and zero-centered *z*-scores.

To ensure that the model is computationally tractable, we computed the rank statistic of each parameter for each inference. The empirical cumulative distributions for μ (black) and σ (red) can be seen in Fig. S8(B). Both distributions appear to be uniform, falling within the 99th percentile of the variation expected from a true uniform distribution. This indicates that the self-consistency relation defined by Eq. (17) holds for this statistical model. With a computationally tractable model in hand, we can now apply the statistical model to our data and verify that data sets drawn from the data-conditioned posterior are indistinguishable from the experimental measurements.

Posterior Predictive Checks. The same statistical model was applied to every unique set of fold-change measurements used in this work. Here, we focus only on the set of fold-change measurements for the double mutant Y20I-Q294V at 50 μ M IPTG. The samples from the posterior distribution conditioned on this dataset can be seen in Fig. S9(A). The joint distribution, shown in the lower left-hand corner, appears fairly symmetric, indicating that μ and σ are independent. There is a slight asymmetry in the sampling of σ , which can be more clearly seen in the corresponding marginal distribution to the right of the joint distribution.

For each MCMC sample of μ and σ , we drew 10 samples from a normal distribution defined by these parameters. From this collection of data sets, we computed the percentiles of the empirical cumulative distribution and plotted them over the data, as can be seen in Fig. S9 (B). We find that the observed data falls within the 99th percentile of the generated data sets. This illustrates that the model can produce data which is identically distributed to the actual experimental measurements, validating our choice of statistical model.

Sensitivity Limits and Systematic Errors in Inference. Considering the results from the prior predictive checks, simulation based calibration, and posterior predictive checks, we can say that the statistical model for inferring μ and σ fold-change from a collection



Fig. S8. Sensitivity measurements and rank statistic distribution of the statistical model estimating μ and σ . (A) Posterior *z*-score of each inference plotted against the posterior shrinkage factor for the parameters μ (red points) and σ (black points). (B) Distribution of rank statistics for μ (red) and σ (black). Gray envelope represents the 99th percentile of a true uniform distribution.

of noisy fold-change measurements is valid and computationally tractable. Upon applying this model to the experimental data of the wild-type strain (where the free energy is theoretically known), we observed that systematic errors arise when the fold-change is exceptionally high or low, making the resulting inference of the free energy inaccurate.

To elucidate the source of this systematic error, we return to a simulation based approach in which the true free energy is known [black points in Fig. S10(A)]. For a range of free energies, we computed the theoretical fold-change prescribed by Eq. (20). For each free energy value, we pulled a value for σ from the prior distribution defined in Eq. (15) and generated a data set of 10 measurements by drawing values from a normal distribution defined by the true fold-change and the drawn value of σ [red points in Fig. S10(A)]. We then sampled the statistical model over these data and inferred the mean fold-change μ [blue points in Fig. S10(A)]. By eye, the inferred points appear to collapse onto the master curve, in many cases overlapping the true values. However, the points with a free energy less than $\approx -2k_BT$ and greater than $\approx 2k_BT$ are slightly above or below the master curve, respectively. This becomes more obvious when the inferred free energy is plotted as a function of the true free energy, shown in Fig. S10(B). Points in which the difference between μ and the neearest boundary (0 or 1) is less than the value of σ are shown as purple or green. When this condition is met, the inferred mean free energy strays from the true value, introducing a systematic error. This suggests that the spread of the fold-change measurements sets the detection limit of fold-change close to either boundary. Thus, the narrower the spread in the fold-change the better the estimate of the fold-change near the boundaries.

These systematic errors can be seen in experimental measurements of the wild-type repressor. Data from Razo-Mejia et al. 2018(2) in which the IPTG titration profiles of seventeen different bacterial strains were measured is shown collapsed onto the master curve in Fig. S10(C) as red points. Here, each point corresponds to a single biological replicate. The inferred mean fold-change μ and 95% credible regions are shown as purple, blue, or green points. The color of these points correspond to the relative value of μ or $1 - \mu$ to σ . The discrepancy between the predicted and inferred free energy of each measurement set can be seen in Fig. S10(D). The significant deviation from the predicted and inferred free energy occurs past the detection limit set by σ . In this work, we therefore opted to not display inferred free energies at the extrema where the inferred fold-change was closer to the boundaries than the correspoding standard deviation, as it reflects limitations in our measurement rather than a deviation from the theoretical predictions.



Fig. S9. MCMC sampling output and posterior predictive checks of the statistical model for the mean fold-change μ and standard deviation σ . (A) Corner plot of sampling output. The joint distribution between σ and μ is shown in the lower left hand corner. Each point is an individual sample. Points are colored by the value of the log posterior with increasing probability corresponding to transitions from blue to yellow. Marginal distributions for each parameter are shown adjacent to the joint distribution. (B) Percentiles of the cumulative distributions from the posterior predictive checks are shown as shaded bars. Data on which the posterior was conditioned are shown as white circles connected by black lines.



Fig. S10. Identification of systematic error in simulated and real data when considering the free energy. (A) The true fold-change (black open circles), simulated fold-change distribution (red points), and inferred mean fold-change (blue) is plotted as a function of the true free energy. Error bars on inferred fold-change correspond to the 95% credible region of the mean fold-change μ . (B) Inferred free energy plotted as a function of the true free energy. Black line indicates perfect agreement between the ground truth free energy and inferred free energy. Blue points correspond to the inferred free energy where the median values of the parameters satisfy the condition $\mu > \sigma$ and $1 - \mu > \sigma$. Purple points correspond to the inferred free energy. Glue points correspond to the inferred free energy. Correspond to those where the inferred mean fold-change $1 - \mu < \sigma$. Error bars correspond to the 95% credible region. (C) Biological replicate data from Razo-Mejia et al. 2018 (2) (red points) plotted as a function of the theoretical free energy. Inferred mean fold-change μ and the 95% credible region are shown as blue points. Purple and green points are colored by the same conditions as in (B). (D) Inferred free energy as a function of the posterior distribution. All inferred values in (A - D) are the median values of the posterior distribution.

4. Additional Characterization of DNA Binding Mutants

In the main text, we estimated the DNA binding energy of each mutant using the mutant strains that had approximately 260 repressors per cell. In this section, we examine the effect of the choice of fit strain on the predictions of both the induction profiles and ΔF for each DNA binding domain mutant.

We applied the statistical model derived in Section 2 for each unique strain of the DNA binding mutants and estimated the DNA binding energy. The median of the posterior distribution along with the upper and lower bounds of the 95% credible region are reported in Table S1. We found that the choice of fitting strain did not strongly influence the estimate of the DNA binding energy. The largest deviations appear for the weakest binding mutants paired with the lowest repressor copy number. In these cases, such as for Q21A, the difference in binding energy between the repressor copy numbers is $\approx 1 k_B T$ which is small compared to the overall DNA binding energy. Using these energies, we computed the predicted induction profiles of each mutant with different repressor copy numbers, shown in Fig. S11. In this plot, the rows correspond to the repressor copy number of the strain used to estimate the DNA binding energy. The columns correspond to the repressor copy number of the strains. The diagonals, shaded in grey, show the induction profile of the fit strain along with the corresponding data. In all cases, we find that the predicted profiles are relatively accurate with the largest deviations resulting from using the lowest repressor copy number as the fit strain.

The predicted change in free energy ΔF using each fit strain can be seen in Fig. S12. In this figure, the rows represent the repressor copy number of the strain to which the DNA binding energy was fit whereas the columns correspond to each mutant. In each plot, we have shown the data for all repressor copy numbers with the fit strain represented by white filled circles. Much as for the induction profiles, we see little difference in the predicted ΔF for each strain, all of which accurately describe the inferred free energies. The ability to accurately predict the majority of the induction profiles of each mutant with repressor copy numbers ranging over two orders of magnitude strengthens our assessment that for these DNA binding domain mutations, only the DNA binding energy is modified.



Fig. S11. Pairwise comparisons of DNA binding mutant induction profiles. Rows correspond to the repressor copy number of the strain used to estimate the DNA binding energy for each mutant. Columns correspond to the repressor copy number of the strains that are predicted. Diagonals in which the data used to estimate the DNA binding energy are shown with a gray background.



Fig. S12. Dependence of fitting strain on ΔF predictions of DNA binding domain mutants. Rows correspond to the repressor copy number used to estimate the DNA binding energy. Columns correspond to the particular mutant. Colored lines are the bounds of the 95% credible region of the predicted ΔF . Open face points indicate the strain to which the DNA binding energy was fit.

Mutant	Repressors	DNA Binding Energy $[k_BT]$
Q21A	60	$-9.8^{+0.2}_{-0.2}$
	124	$-10.3\substack{+0.1\\-0.1}$
	260	$-11.0\substack{+0.1\\-0.1}$
	1220	$-11.3\substack{+0.1\\-0.1}$
Q21M	60	$-15.83\substack{+0.08\\-0.08}$
	124	$-15.7\substack{+0.1\\-0.1}$
	260	$-15.43\substack{+0.07\\-0.06}$
	1220	$-15.27\substack{+0.07\\-0.07}$
Y20I	60	$-9.4\substack{+0.3\\-0.3}$
	124	$-9.5^{+0.1}_{-0.1}$
	260	$-9.9\substack{+0.1\\-0.1}$
	1220	$-10.1\substack{+0.2\\-0.2}$

Table S1. Estimated DNA binding energy for DNA binding domain mutants with different repressor copy numbers. Median of the posterior distribution with the upper and lower bounds of the 95% credible region are reported.

5. Bayesian Parameter Estimation for Inducer Binding Domain Mutants

In the main text, we put forward two naïve hypotheses for which parameters of Eq. (7) are affected by mutations in the inducer binding domain of the repressor. The first hypothesis was that only the inducer dissociation constants, K_A and K_I , were perturbed from their wild-type values. Another hypothesis was that the inducer dissociation constants were affected in addition to the energetic difference between the active and inactive states of the repressor, $\Delta \varepsilon_{AI}$.

In this section, we first derive the statistical model for each hypothesis and then perform a series of diagnostic tests that expose the inferential limitations of each model. With well calibrated statistical models, we then apply each to an induction profile of the inducer binding mutant Q294K and assess the validity of each hypothesis. To understand the statistical models for each hypothesis, only the subsection *Building A Generative Statistical Model* is necessary.

Building a Generative Statistical Model. For both hypotheses, we assume that the underlying physical model defined in Eq. (7) is the same while a subset of the parameters are modified. As the fold-change measurements for each biological replicate are statistically independent, we can assume that they are normally distributed about the theoretical fold-change value. Thus, for each model, we must include a parameter σ which is the standard deviation of the distribution of fold-change measurements. For the first hypothesis, in which only K_A and K_I are changed, we are interested in sampling the posterior distribution

$$g(K_A, K_I, \sigma \mid y) \propto f(y \mid K_A, K_I, \sigma)g(K_A)g(K_I)g(\sigma),$$
^[27]

where *y* corresponds to the set of fold-change measurements. In the above model, we have assumed that the priors for K_A and K_I are independent. It is possible that it is more appropriate to assume that they are dependent and that a single prior distribution captures both parameters, $g(K_A, K_I)$. However, assigning this prior is more difficult and requires strong knowledge *a priori* about the relationship between them. Therefore, we continue under the assumption that the priors are independent.

The generic posterior given in Eq. (27) can be extended to evaluate the second hypothesis in which $\Delta \varepsilon_{AI}$ is also modified,

$$g(K_A, K_I, \Delta \varepsilon_{AI}, \sigma \mid y) \propto f(y \mid K_A, K_I, \Delta \varepsilon_{AI}, \sigma) g(K_A) g(K_I) g(\Delta \varepsilon_{AI}) g(\sigma)$$
^[28]

where we have included $\Delta \varepsilon_{AI}$ as an estimated parameter and assigned a prior distribution.

As we have assumed that the fold-change measurements across replicates are independent and normally distributed, the likelihoods for each hypothesis can be written as

$$f(y \mid K_A, K_I, \sigma) \sim \text{Normal}\{\mu(K_A, K_I), \sigma\},$$
[29]

for the first hypothesis and

$$f(y | K_A, K_I, \Delta \varepsilon_{AI}, \sigma) \sim \text{Normal}\{\mu(K_A, K_I, \Delta \varepsilon_{AI}), \sigma\},$$
[30]

for the second. Here, we have assigned $\mu(...)$ as the mean of the normal distribution as a function of the parameters defined by our fold-change equation, Eq. (7).

With a likelihood distribution in hand, we now turn toward assigning functional forms to each prior distribution. As we have used in the previous sections [Sec. 2 and Sec. 3], we can assign a half-normal prior for σ with a standard deviation of 0.1, namely,

$$g(\sigma) \sim \text{HalfNormal}\{0, 0.1\}.$$
 [31]

It is important to note that the inducer dissociation constants K_A and K_I are scale invariant, meaning that a change from 0.1 μ M to 1 μ M yields a decrease in affinity equal to a change from 10 μ M to 100 μ M. As such, it is better to sample the dissociation constants on a logarithmic scale. We can assign a log normal prior for each dissociation constant as

$$g(K_A) = \frac{1}{K_A \sqrt{2\pi\phi^2}} \exp\left[-\frac{(\log\frac{K_A}{1\,\mu\mathrm{M}} - \mu_{K_A})^2}{2\phi^2}\right],$$
[32]

or with the short-hand notion of

$$g(K_A) \sim \text{LogNormal}\{\mu_{K_A}, \phi\}$$
 [33]

For K_A , we assigned a mean $\mu_{K_A} = 2$ and a standard deviation $\phi = 2$. For K_I , we chose a mean of $\mu_{K_I} = 0$ and $\phi = 2$, capturing our prior knowledge that $K_A > K_I$ for the wild-type LacI. While the prior distributions are centered differently, they both show extensive overlap, permitting mutations in which $K_A < K_I$. For $\Delta \varepsilon_{AI}$, we assign a normal distribution of the prior centered at 0 with a standard deviation of $5 k_B T$,

$$g(\Delta \varepsilon_{AI}) \sim \text{Normal}\{0, 5\}.$$
 [34]

This permits values of $\Delta \varepsilon_{AI}$ that are above or below zero, meaning that the inactive state of the repressor can be either more or less energetically favorable to the active state. A standard deviation of $5k_BT$ permits a wide range of energies with $+5k_BT$ and $-5k_BT$ corresponding to $\approx 99.5\%$ and $\approx 0.5\%$ of the repressors being active in the absence of inducer, respectively.

20 of 37 Griffin Chure, Manuel Razo-Mejia, Nathan M. Belliveau, Tal Einav, Stephanie L. Barnes, Mitchell Lewis, Rob Phillips

Prior predictive checks. To ensure that these choices of prior distributions are appropriate, we performed prior predictive checks for each hypothesis as previously described in Section 2. We drew 1000 values from the prior distributions shown in Fig. S13(A) for K_A , K_I , and $\Delta \epsilon_{AI}$. Using the draws from the K_A , and K_I priors alone, we generated datasets of \approx 70 measurements. The percentiles of the fold-change values drawn for the 1000 simulations is shown in the top panel of Fig. S13(B).

It can be seen that in the absence of inducer, the fold-change values are close to zero and are with distributed about the leakiness value due to σ . This is in contrast to the data sets generated when $\Delta \varepsilon_{AI}$ is permitted to vary along with K_A and K_I . In the bottom panel of Fig. S13(B), the fold-change when c = 0 can extend above 1.0 which is possible only when $\Delta \varepsilon_{AI}$ is included, which sets what fraction of the repressors is active. Under both hypotheses, the 99th percentile of the fold-change extends to just above 1 or just below 0, which matches our intuition of how the data should behave. Given these results, we are satisfied with these choices of priors and continue onto the next level of calibration of our model.



Fig. S13. Prior predictive checks for two hypotheses of inducer binding domain mutants. (A) Probability density functions for K_A , K_I , $\Delta \varepsilon_{AI}$, and σ . Black points correspond to draws from the distributions used for prior predictive checks. (B) Percentiles of the simulated data sets using draws from the K_A and K_I distributions only (top, red bands) and using draws from K_A , K_I , and $\Delta \varepsilon_{AI}$ (bottom, blue bands).

Simulation Based Calibration. With an appropriate choice of priors, we turn to simulation based calibration to root out any pathologies lurking in the model itself or the implementation through MCMC. For each parameter under each model, we compute the *z*-score and shrinkage of each inference, shown in Fig. S14. Under the first hypothesis in which K_A and K_I are the only perturbed parameters [Fig. S14(A)], we see all parameters have *z*–scores clustered around 0, indicating that the value of the ground-truth is being accurately estimated through the inference. While the shrinkage for σ is close to 1 (indicating the prior is being informed by the data), the shrinkage for K_A and K_I is heavily tailed with some values approaching zero. This is true for both statistical models, indicating that for some values of K_A and K_I , the parameters are difficult to pin down with high certainty. In the application of these models to data, this will be revealed as large credible regions in the reported parameters. Under the second hypothesis in which all allosteric parameters are allowed to change, we see moderate shrinkage for $\Delta \varepsilon_{AI}$ [purple points in Fig. S14(B)] with the minimum shrinkage being around 0.5. The samples resulting in low shrinkage correspond to values of $\Delta \varepsilon_{AI}$ that are highly positive or highly negative, in which small changes in the active fraction of repressors cannot be accurately measured through our model. However, the median shrinkage for $\Delta \varepsilon_{AI}$ is approximately 0.92, meaning that the the data highly informed the prior distributions for the majority of the inferences. The rank distributions for all parameters under each model appear to be highly uniform, indicating that both statistical models are computationally tractable.

With knowledge of the caveats of estimating K_A and K_I for both models, we proceed with our analysis and examine how accurately these models can capture the phenomenology of the data.

Posterior Predictive Checks. With a properly calibrated statistical model for each hypothesis, we now apply it to a representative dataset. While each model was applied to each inducer binding domain mutant, we only show the application to the mutant Q294K with 260 repressors per cell paired with the native *lac* operator O2.



Fig. S14. Simulation based calibration of statistical models for inducer binding domain mutants. (A) Sensitivity statistics and rank distribution for a statistical model in which K_A and K_I are the only parameters permitted to vary. (B) Sensitivity statistics and rank distribution for a model in which all allosteric parameters K_A , K_I , and $\Delta \varepsilon_{AI}$ are allowed to be modified by the mutation. Gray envelope in the bottom plots correspond to the 99th percentile of variation expected from a true uniform distribution.

The results from applying the statistical model in which only K_A and K_I can change is shown in Fig. S15. The joint and marginal distributions for each parameter [Fig. S15(A)] reveal a strong correlation between K_A and K_I whereas all other parameters are symmetric and independent. While the joint and marginal distributions look well behaved, the percentiles of the posterior predictive checks [Fig. S15(B)] are more suspect. While all data falls within the 95th percentile, the overall trend of the data is not well predicted. Furthermore, the percentiles expand far below zero, indicating that the sampling of σ is compensating for the leakiness in the data being larger than it should be if only K_A and K_I were the changing parameters.

We see significant improvement when $\Delta \varepsilon_{AI}$ is permitted to vary in addition to K_A and K_I . Fig. S16(A) shows the joint and marginal distributions between all parameters from the MCMC sampling. We still see correlation between K_A and K_I , although it is not as strong as in the case where they are the only parameters allowed to change due to the mutation. We also see that the marginal distribution for σ has shrunk significantly compared to the marginal distribution in Fig. S15(A). The percentiles of the posterior predictive checks, shown in Fig. S16(B) are much more in line with the experimental measurements, with the 5th percentile following the data for the entire induction profile.

In this section we have presented two hypotheses for the minimal parameter set needed to describe the inducer binding mutations, derived a statistical model for each, thoroughly calibrated its behavior, and applied it to a representative data set. The posterior predictive checks [Fig. S15 and Fig. S16] help us understand which hypothesis is more appropriate for that particular mutant. The incredibly wide percentiles and significant change in the leakiness that result from a model in which only K_A and K_I are perturbed suggests that more than those two parameters should be changing. We see significant improvement in the description of the data when $\Delta \varepsilon_{AI}$ is altered, indicating that it is the more appropriate hypothesis of the two.



Fig. S15. Posterior predictive checks for inducer binding domain mutants where only K_A and K_I are changed. (A) MCMC sampling output for each parameter. Joint distributions are colored by the value of the log posterior with increasing probability corresponding to transition from blue to yellow. (B) Percentiles of the data generated from the likelihood distribution for each sample of K_A , K_I , and σ . Overlaid points are the experimentally observed measurements.

(A)



(B)

Fig. S16. Posterior predictive checks for inducer binding domain mutants where all allosteric parameters can change. (A) MCMC sampling output for all parameters. Joint distributions are colored by the value of the log posterior with increasing probability corresponding to the transition from blue to yellow. Marginal distributions are shown adjacent to each joint distribution. (B) Percentiles of the data generated from the likelihood for each sample of K_A , K_I , $\Delta \varepsilon_{AI}$, and σ . The corresponding experimental data for Q294K are shown as black open-faced circles.

(A)

6. Additional Characterization of Inducer Binding Domain Mutants

To predict the induction profiles of the inducer binding mutants, we used only the induction profile of each mutant paired with the native O2 *lac* operator to infer the parameters. Here, we examine the influence the choice of fit strain has on the predictions of the induction profiles and ΔF for each mutant.

In the main text, we dismissed the hypothesis that only K_A and K_I were changing due to the mutation and based the fit to a single induction profile. In Fig. S17, the fits and predictions for each mutant paired with each operator sequence queried. Here, the rows correspond to the operator sequence of the fit strain while the columns correspond to the operator sequence of the predicted strain. The diagonals, colored in gray, show the fit induction profiles and the corresponding data. Regardless of the choice of fit strain, the predicted induction profiles of the repressor paired with the O3 operator are poor, with the leakiness in each case being significantly underestimated. We also see that fitting to O3 results in poor predictions with incredibly wide credible regions for the other two operators. In Razo-Mejia et al. 2018 (2), we also found that fitting K_A and K_I to the induction profile of O3 generally resulted in poor predictions of the other strains with comparably wide credible regions.

When $\Delta \varepsilon_{AI}$ is included as a parameter, however, the predictive power is improved for all three operators, as can be seen in Fig. **S18**. While the credible regions are still wide when fit to the O3 operator, they are much narrower than under the first hypothesis. We emphasize that we are able to accurately predict the leakiness of nearly every strain by redetermining $\Delta \varepsilon_{AI}$ whereas the leakiness was not predicted when only K_A and K_I were considered. Thus, we conclude that all three allosteric parameters K_A , K_I , and $\Delta \varepsilon_{AI}$ are modified for these four inducer binding domain mutations. The values of the inferred parameters are reported in Table S2.

We also examined the effect the choice of fit strain has on the predicted ΔF , shown in Fig. S19. We find that the predictions agree with the data regardless of the choice of fit strain. One exception is the prediction of the Q294K ΔF when the parameters fit to the O3 induction profile are used. As the induction profile for Q294K paired with O3 is effectively flat at a fold-change of 1, it is difficult to properly estimate the parameters of our sigmoidal function. We note all measurements of ΔF for Q294K are described by using either the parameters fit to either O1 or O3 induction profiles, suggesting that the choice of fit strain makes little difference.



Fig. S17. Pairwise comparison of fit strain versus predictions assuming only K_A and K_I are influenced by the mutation. Rows correspond to the operator sequence of the strain used for the parameter inference. Columns correspond to the operator sequence of the predicted strain. Colors identify the mutation. Diagonal positions (gray background) show the induction fit strain and profiles.



Fig. S18. Pairwise comparison of fit strain versus predictions assuming all allosteric parameters are affected by the mutation. Rows correspond to the operator of the strain used to fit the parameters. Columns correspond to the operator of the strains whose induction profile is predicted. Mutants are identified by color. Diagonals (gray background) show the induction profiles of the strain to which the parameters were fit.



Fig. S19. Comparison of choice of fit strain on predicted ΔF profiles. Rows correspond to the operator of the strain to which the parameters were fit. Columns correspond to mutations. Points are colored by their operator sequence. The data corresponding to the operator of the fit strain are shown as white-faced points.

Mutant	Operator	ΚA [μ M]	Κ Ι [μ Μ]	$\Delta \varepsilon_{AI}$ [k _b T]
F164T	O1	290^{+60}_{-56}	$1^{+4}_{-0.98}$	4^{+5}_{-3}
	O2	165^{+90}_{-65}	3^{+6}_{-3}	1^{+5}_{-2}
	O3	$110\substack{+700 \\ -105}$	7^{+5}_{-4}	$-0.9\substack{+0.4 \\ -0.3}$
Q294K	01	> 1 mM	$410\substack{+150 \\ -100}$	$-3.2\substack{+0.1\\-0.1}$
	O2	> 1 mM	$310\substack{+70 \\ -60}$	$-3.11\substack{+0.07\\-0.07}$
	O3	$10\substack{+200 \\ -10}$	1^{+9}_{-1}	-7^{+3}_{-5}
Q294R	01	3^{+27}_{-3}	2^{+20}_{-2}	$-1.9\substack{+0.4 \\ -0.3}$
	O2	9^{+20}_{-9}	8^{+20}_{-8}	$-2.32\substack{+0.01\\-0.09}$
	O3	6^{+24}_{-6}	9^{+30}_{-9}	$-2.6\substack{+0.4\\-0.5}$
Q294V	01	> 1 mM	3^{+13}_{-3}	6^{+4}_{-4}
	O2	$650\substack{+450 \\ -250}$	8^{+8}_{-8}	3^{+6}_{-3}
	O3	$100\substack{+400 \\ -90}$	22^{+33}_{-18}	$0.1\substack{+0.8 \\ -0.6}$

Table S2. Inferred values of K_A , K_I , and $\Delta \varepsilon_{AI}$ for inducer binding domain mutants. Values reported are the mean of the posterior distribution with the upper and lower bounds of the 95% credible region.

7. Parameter Estimation Using All Induction Profiles

In the main text and Sec. 4 and 6 of this supplementary text, we have laid out our strategy for inferring the the various parameters of our model to a single induction profile and using the resulting values to predict the free energy and induction profiles of other strains. In this section, we estimate the parameters using all induction profiles of a single mutant and using the estimated values to predict the free energy profiles.

The inferred DNA binding energies considering induction profiles of all repressor copy numbers for the three DNA binding mutants are reported in Tab. S3. These parameters are close to those reported in Tab. S1 for each repressor copy number with Q21A showing the largest differences. The resulting induction profiles and predicted change in free energy for these mutants can be seen in Fig. ??. Overall, the induction profiles match the data to an appreciable agree. We acknowledge that even when using *all* repressor copy numbers, the fit to Q21A remains imperfect. However we contend that this disagreement is comparable to that observed in (2) which described the induction profile of the wild-type repressor. We find that the predicted change in free energy [bottom row in Fig. ??(B)] narrows compared to that in Fig. S12 and Fig. 3 of the main text, confirming that considering all induction profiles improves our inference of the most-likely DNA binding energy. There appears to be a very slight trend in the ΔF for Q21A at higher inducer concentrations, though the overall change in free energy from 0 to 5000 μ M IPTG is small.

We also estimated the allosteric parameters (K_A , K_I , and $\Delta \varepsilon_{AI}$) for all inducer binding domain mutations using the induction profiles of all three operator sequences. The values, reported in Tab. S4 are very similar to those estimated from a single induction profile (Tab. S2). We note that for Q294R, it is difficult to properly estimate the values for K_A and K_I as the observed induction profile is approximately flat. The induction profiles and predicted change in free energy for each inducer binding mutant is shown in Fig. S21. We see notable improvement in the agreement between the induction profiles and the observed data, indicating that considering all data significantly shrinks the uncertainty of each parameter. The predicted change in free energy is also improved compared to that shown in Fig. S19. We emphasize that the observed free energy difference for each point assumes no knowledge of the underlying parameters and comes directly from measurements. The remarkable agreement between the predicted free energy and the observations illustrates that redetermining the allosteric parameters is sufficient to describe how the free energy changes as a result of the mutation.

Table S3. Estimated DNA binding energies for each DNA binding domain mutant using all repressor copy numbers

Mutant	$\Delta \varepsilon_{RA} [k_B T]$
Y20I	$-9.81\substack{+0.04 \\ -0.08}$
Q21A	$-10.60\substack{+0.07\\-0.07}$
Q21M	$-15.61\substack{+0.05\\-0.05}$



Fig. S20. Induction profiles and predicted change in free energy using parameters estimated from the complete data sets. Top row shows fold-change measurements (points) as mean and standard error with ten to fifteen biological replicates. Shaded lines correspond to the 95% credible regions of the induction profiles using the estimated values of the DNA binding energies reported in Tab. S3. Bottom row shows the 95% credible regions of the predicted change in free energy (shaded lines) along with the inferred free energy of data shown in the top row. In all plots, the inducer concentration is shown on a symmetric log scale with linear scaling between 0 and $10^{-2} \,\mu$ M and log scaling elsewhere.

Table S4. Estimated values for K_A , K_I , and $\Delta \varepsilon_{AI}$ for inducer binding domain mutations using induction profiles of all operator sequences.

Mutant	$K_A \ [\mu M]$	$K_I \ [\mu M]$	$\Delta \varepsilon_{AI} [k_B T]$
F164T	300^{+60}_{-60}	$12.7\substack{+0.1 \\ -0.1}$	$-0.9\substack{+0.3 \\ -0.3}$
Q294K	$> 1 \ \mathrm{mM}$	330^{+60}_{-70}	$-3.17\substack{+0.07 \\ -0.07}$
Q294R	$> 1 \ \mathrm{mM}$	$> 1 \mathrm{~mM}$	$-2.4\substack{+0.2\\-0.2}$
Q294V	$> 1 \ \mathrm{mM}$	53^{+17}_{-13}	$0^{+0.3}_{-0.3}$



Fig. S21. Induction profiles and predicted change in free energy using parameters estimated from the complete data sets for inducer binding domain mutnats. Top row shows fold-change measurements (points) as mean and standard error with ten to fifteen biological replicates. Shaded lines correspond to the 95% credible regions of the induction profiles using the estimated values of the allosteric parameters reported in Tab. S4. Bottom row shows the 95% credible regions of the predicted change in free energy (shaded lines) along with the inferred free energy of data shown in the top row. In all plots, the inducer concentration is shown on a symmetric log scale with linear scaling between 0 and $10^{-2} \mu M$ and log scaling elsewhere.

Strain and Oligonucleotide Information

Table S5. Escherichia coli strains used in this work

Class	Lacl Mutant	Operator	Rep. per Cell	Genotype	Plasmid
-	-	-	22	MG1655::∆ <i>lacZYA</i>	pZS4*1-mCherry
-	-	01	0	MG1655::∆ <i>laclZYA;galK</i> <>25O1+11-YFP	pZS4*1-mCherry
-	-	O2	0	MG1655::∆ <i>laclZYA;galK</i> <>25O2+11-YFP	pZS4*1-mCherry
-	-	O3	0	MG1655::∆ <i>laclZYA;galK</i> <>25O3+11-YFP	pZS4*1-mCherry
WT	WT	O2	260	MG1655::∆ <i>laclZYA;galK<</i> >25O2+11-YFP; <i>ybcN<</i> >3*1-RBS1027Lacl	pZS4*1-mCherry
DNA	Y20I	O2	60	MG1655::∆ <i>laclZYA;galK</i> <>25O2+11-YFP; <i>ybcN</i> <>3*1-RBS1147Lacl(Y20I)	pZS4*1-mCherry
DNA	Y20I	O2	124	MG1655::∆ <i>laclZYA;galK<</i> >25O2+11-YFP; <i>ybcN<</i> >3*1-RBS446Lacl(Y20I)	pZS4*1-mCherry
DNA	Y20I	O2	260	MG1655::∆ <i>laclZYA;galK<</i> >25O2+11-YFP; <i>ybcN<</i> >3*1-RBS1027Lacl(Y20I)	pZS4*1-mCherry
DNA	Y20I	O2	1220	MG1655::∆ <i>laclZYA;galK<</i> >25O2+11-YFP; <i>ybcN<</i> >3*1-RBS1Lacl(Y20I)	pZS4*1-mCherry
DNA	Q21A	O2	60	MG1655::∆ <i>laclZYA;galK<</i> >25O2+11-YFP; <i>ybcN<</i> >3*1-RBS1147Lacl(Q21A)	pZS4*1-mCherry
DNA	Q21A	O2	124	MG1655::∆ <i>laclZYA;galK<</i> >25O2+11-YFP; <i>ybcN<</i> >3*1-RBS446Lacl(Q21A)	pZS4*1-mCherry
DNA	Q21A	O2	260	MG1655::∆ <i>laclZYA;galK<</i> >25O2+11-YFP; <i>ybcN<</i> >3*1-RBS1027Lacl(Q21A)	pZS4*1-mCherry
DNA	Q21A	O2	1220	MG1655::∆ <i>laclZYA;galK</i> <>25O2+11-YFP; <i>ybcN</i> <>3*1-RBS1Lacl(Q21A)	pZS4*1-mCherry
DNA	Q21M	O2	60	MG1655::∆ <i>laclZYA;galK</i> <>25O2+11-YFP; <i>ybcN</i> <>3*1-RBS1147Lacl(Q21M)	pZS4*1-mCherry
DNA	Q21M	O2	124	MG1655::∆ <i>laclZYA;galK</i> <>25O2+11-YFP; <i>ybcN</i> <>3*1-RBS446Lacl(Q21M)	pZS4*1-mCherry
DNA	Q21M	O2	260	MG1655::∆ <i>laclZYA;galK</i> <>25O2+11-YFP; <i>ybcN</i> <>3*1-RBS1027Lacl(Q21M)	pZS4*1-mCherry
DNA	Q21M	O2	1220	MG1655::∆ <i>laclZYA;galK</i> <>25O2+11-YFP; <i>ybcN</i> <>3*1-RBS1Lacl(Q21M)	pZS4*1-mCherry
IND	F164T	01	260	MG1655::∆ <i>laclZYA;galK</i> <>25O1+11-YFP; <i>ybcN</i> <>3*1-RBS1027Lacl(F164T)	pZS4*1-mCherry
IND	F164T	O2	260	MG1655::∆ <i>laclZYA;galK</i> <>25O2+11-YFP; <i>ybcN</i> <>3*1-RBS1027Lacl(F164T)	pZS4*1-mCherry
IND	F164T	O3	260	MG1655::∆ <i>laclZYA;galK</i> <>25O3+11-YFP; <i>ybcN</i> <>3*1-RBS1027Lacl(F164T)	pZS4*1-mCherry
IND	Q294V	01	260	MG1655::∆ <i>laclZYA;galK</i> <>25O1+11-YFP; <i>ybcN</i> <>3*1-RBS1027Lacl(Q294V)	pZS4*1-mCherry
IND	Q294V	O2	260	MG1655::∆ <i>laclZYA;galK<</i> >25O2+11-YFP; <i>ybcN<</i> >3*1-RBS1027Lacl(Q294V)	pZS4*1-mCherry
IND	Q294V	O3	260	MG1655::∆ <i>laclZYA;galK</i> <>25O3+11-YFP; <i>ybcN</i> <>3*1-RBS1027Lacl(Q294V)	pZS4*1-mCherry
IND	Q294K	01	260	MG1655::∆ <i>laclZYA;galK</i> <>25O1+11-YFP; <i>ybcN</i> <>3*1-RBS1027Lacl(Q294K)	pZS4*1-mCherry
IND	Q294K	O2	260	MG1655::∆ <i>laclZYA;galK<</i> >25O2+11-YFP; <i>ybcN<</i> >3*1-RBS1027Lacl(Q294K)	pZS4*1-mCherry
IND	Q294K	O3	260	MG1655::∆ <i>laclZYA;galK</i> <>25O3+11-YFP; <i>ybcN</i> <>3*1-RBS1027Lacl(Q294K)	pZS4*1-mCherry
IND	Q294R	01	260	MG1655::∆ <i>laclZYA;galK</i> <>25O1+11-YFP; <i>ybcN</i> <>3*1-RBS1027Lacl(Q294R)	pZS4*1-mCherry
IND	Q294R	O2	260	MG1655::∆ <i>laclZYA;galK</i> <>25O2+11-YFP; <i>ybcN</i> <>3*1-RBS1027Lacl(Q294R)	pZS4*1-mCherry
IND	Q294R	O3	260	MG1655::∆ <i>laclZYA;galK</i> <>25O3+11-YFP; <i>ybcN</i> <>3*1-RBS1027Lacl(Q294R)	pZS4*1-mCherry
DBL	Y20I-F164T	O2	260	MG1655::∆ <i>laclZYA;galK</i> <>25O2+11-YFP; <i>ybcN</i> <>3*1-RBS1027Lacl(Y20IF164T)	pZS4*1-mCherry
DBL	Y20I-Q294V	O2	260	MG1655::∆ <i>laclZYA;galK</i> <>25O2+11-YFP; <i>ybcN</i> <>3*1-RBS1027Lacl(Y20IQ294V)	pZS4*1-mCherry
DBL	Y20I-Q294K	O2	260	MG1655::∆ <i>laclZYA;galK</i> <>25O2+11-YFP; <i>ybcN</i> <>3*1-RBS1027Lacl(Y20IQ294K)	pZS4*1-mCherry
DBL	Q21A-F164T	O2	260	MG1655::∆ <i>laclZYA;galK</i> <>25O2+11-YFP; <i>ybcN</i> <>3*1-RBS1027Lacl(Q21AF164T)	pZS4*1-mCherry
DBL	Q21A-Q294V	O2	260	MG1655::∆ <i>laclZYA;galK</i> <>25O2+11-YFP; <i>ybcN</i> <>3*1-RBS1027Lacl(Q21AQ294V)	pZS4*1-mCherry
DBL	Q21A-Q294K	O2	260	MG1655::∆ <i>laclZYA;galK</i> <>25O2+11-YFP; <i>ybcN</i> <>3*1-RBS1027Lacl(Q21AQ294K)	pZS4*1-mCherry
DBL	Q21M-F164T	O2	260	MG1655::∆ <i>laclZYA;galK<></i> 25O2+11-YFP; <i>ybcN<></i> 3*1-RBS1027Lacl(Q21MF164T)	pZS4*1-mCherry
DBL	Q21M-Q294V	O2	260	MG1655::∆ <i>laclZYA;galK</i> <>25O2+11-YFP; <i>ybc</i> N<>3*1-RBS1027Lacl(Q21MQ294V)	pZS4*1-mCherry
DBL	Q21M-Q294K	O2	260	MG1655::\[]/aclZYA;galK<>25O2+11-YFP;ybcN<>3*1-RBS1027Lacl(Q21MQ294K)	pZS4*1-mCherry

Table S6. Oligonucleotides used for mutant generation.

Primer Name	Sequence (5'→3')	Description	Method
10.1	${\tt acctctgcggagggaagcgtgaacctctcacaagacggcatcaaattacactagcaacaccagaacagc}$	Integration into ybcN locus	λ -Red Recombineering
10.3	${\tt ctgtagatgtgtccgttcatgacacgaataagcggtgtagccattacgccggctaatgcacccagtaagg}$	Integration into ybcN locus	λ -Red Recombineering
GCMWC-001	ccggcatactctgcgaca	Amplification of plasmid	QuickChange Mutagenesis
GCMWC-002	gtgtctcttat <u>ATG</u> accgtttcccgc	Q21M Mutation (CAG \rightarrow ATG)	QuickChange Mutagenesis
GCMWC-003	tgtctcttat <u>GCG</u> accgtttcccgc	Q21A Mutation (CAG \rightarrow GCG)	QuickChange Mutagenesis
GCMWC-004	gttaacggcgggatataac	Amplification of plasmid	QuickChange Mutagenesis
GCMWC-005	caccatcaaa <u>GTG</u> gattttcgcctgc	Q294V Mutation (CAG $ ightarrow$ GTG)	QuickChange Mutagenesis
GCMWC-006	caccatcaaa <u>AAG</u> gattttcgcc	Q294K Mutation (CAG \rightarrow AAG)	QuickChange Mutagenesis
GCMWC-007	cagtattatt <u>ACC</u> tcccatgaagacgg	F164T Mutation (TTC \rightarrow ACC)	QuickChange Mutagenesis
GCMWC-008	ttgatgggtgtctggtcag	Amplification of plasmid	QuickChange Mutagenesis
GCMWC-009	gcatactctgcgacatcgtataa	Amplification of plasmid	QuickChange Mutagenesis
GCMWC-010	cggtgtctct <u>ATT</u> cagaccgtttc	Y20I Mutation (TAT \rightarrow ATT)	QuickChange Mutagenesis
GCMWC-017	ccatcaaa <u>AGG</u> gattttcgcctgctggggcaaaccag	Q294R Mutation (CAG \rightarrow AGG)	Gibson Assembly
GCMWC-018	ggcgaaaatc <u>CCT</u> tttgatggtggttaacggcggg	Q294R Mutation (CTG \rightarrow CCT)	Gibson Assembly

References

- 1. Garcia HG, Phillips R (2011) Quantitative dissection of the simple repression input-output function. *Proceedings of the National Academy of Sciences of the United States of America* 108(29):12173–8.
- 2. Razo-Mejia M, et al. (2018) Tuning Transcriptional Regulation through Signaling: A Predictive Theory of Allosteric Induction. *Cell Systems* 6(4):456–469.
- 3. Ackers GK, Johnson AD, Shea MA (1982) Quantitative model for gene regulation by lambda phage repressor. *Proceedings of the National Academy of Sciences of the United States of America* 79(4):1129–33.
- 4. Buchler NE, Gerland U, Hwa T (2003) On schemes of combinatorial transcription logic. *Proceedings of the National Academy of Sciences of the United States of America* 100(9):5136–41.
- 5. Vilar JM, Leibler S (2003) DNA Looping and Physical Constraints on Transcription Regulation. *Journal of Molecular Biology* 331(5):981–989.
- 6. Bintu L, et al. (2005) Transcriptional regulation by the numbers: Applications. *Current Opinion in Genetics and Development* 15(2):125–135.
- 7. Kuhlman T, Zhang Z, Saier MH, Hwa T (2007) Combinatorial transcriptional control of the lactose operon of *Escherichia coli*. *Proceedings of the National Academy of Sciences of the United States of America* 104(14):6043–6048.
- 8. Daber R, Sochor MA, Lewis M (2011) Thermodynamic analysis of mutant lac repressors. *Journal of Molecular Biology* 409(1):76–87.
- 9. Brewster RC, et al. (2014) The transcription factor titration effect dictates level of gene expression. Cell 156(6):1312–1323.
- Weinert FM, Brewster RC, Rydenfelt M, Phillips R, Kegel WK (2014) Scaling of Gene Expression with Transcription-Factor Fugacity. *Physical Review Letters* 113(25):1–5.
- 11. Klumpp S, Hwa T (2008) Growth-rate-dependent partitioning of RNA polymerases in bacteria. *Proceedings of the National Academy of Sciences of the United States of America* 105(51):20245–20250.
- 12. Brewster RC, Jones DL, Phillips R (2012) Tuning Promoter Strength through RNA Polymerase Binding Site Design in *Escherichia coli. PLoS Computational Biology* 8(12).
- 13. Schmidt A, et al. (2015) The quantitative and condition-dependent Escherichia coli proteome. Nature Biotechnology 34(1):104–110.
- 14. Monod J, Wyman J, Changeux JP (1965) On the Nature of Allosteric Transitions: A Plausible Model. *Journal of molecular biology* 12:88–118.
- 15. Phillips R, Milo R (2015) Rates and duration in Cell Biology by the Numbers.
- 16. Milo R, Jorgensen P, Moran U, Weber G, Springer M (2010) BioNumbers—the database of key numbers in molecular and cell biology. *Nucleic Acids Research* 38(suppl_1):D750–D753.
- 17. Good IJ, 1950 (1950) Probability and the Weighing of Evidence . (New York City).
- 18. Talts S, Betancourt M, Simpson D, Vehtari A, Gelman A (2018) Validating Bayesian Inference Algorithms with Simulation-Based Calibration. *arXiv:1804.06788 [stat]*.
- 19. Carpenter B, et al. (2017) Stan: A probabilistic programming language. Journal of Statistical Software, Articles 76(1):1–32.