Supporting Information for

"Resolving molecular heterogeneity with single-molecule centrifugation"

Yi Luo‡, Jeffrey Chang‡, Darren Yang‡, J Shepard Bryan IV, Molly MacIsaac, Steve Pressé, Wesley P. Wong* (‡equal contribution; *email: wesley.wong@childrens.harvard.edu)

Table of Contents

List of Figures

Section S1 Methods

Conjugating Antibodies to DNA Oligos

In order to attach antibodies to the nanoswitch construct, IgG antibody proteins of interest were first coupled to DNA oligos complementary to a specific region on the nanoswitch backbone. The antibodies tested in this work were monoclonal anti-fluorescein (Invitrogen 31242), monoclonal anti-digoxigenin (Invitrogen 700772), and polyclonal anti-fluorescein (Invitrogen A-889). Before conjugation, the antibodies were washed 3 times using Zeba desalting columns (7K MWCO, 0.5mL, Thermo Scientific 89882) with PBS buffer pH8.0 to remove free amine and azide in the storage buffer. Antibody concentrations before and after buffer exchange were determined via OD280 readings. The post buffer exchange concentration was adjusted to 0.5mg/ml (3.3uM) by dilution with PBS. The antibody was then activated by converting some surface amine groups into dibenzocyclooctyne (DBCO) groups for copper-free click chemistry using the short linker DBCO-PEG4-NHS (Sigma 764019). The DBCO-PEG4-NHS was first dissolved in DMSO (anhydrous, Invitrogen D12345) to 0.4mM, and then spiked into the antibody in PBS solution, at ratios ranging from 1x to 3x molar excess of DBCO-PEG4-NHS to antibody. The reaction vial was flushed with argon gas briefly, sealed with Parafilm, and incubated at room temperature for 1 hour with rotation.

After incubation, the reaction mixture was washed 3 times again with Zeba columns to remove excess DBCO-PEG4-NHS and transferred into PBS buffer at a pH of 7.4. DNA oligos with a 3' azide modification were then added at a 10x molar ratio to the DBCO-activated antibody. The reaction mixture was incubated at room temperature for 2 hours with rotation, followed by size exclusion FPLC using a Superdex 200 Increase 10/300 GL column (Cytiva, 28-9909-44). Elusion fractions were resolved by native TBE gel electrophoresis. The fractions with pure antibody-DNA conjugates were pooled. The purified conjugates were mixed with glycerol (Sigma G5516) to 40% (v/v) and stored at -20°C.

DNA Nanoswitch Preparation

The DNA nanoswitch construct was prepared based on protocols previously published by our team.[1,2] Briefly, a set of complementary DNA oligos (Integrated DNA Technologies)[3,4] were hybridized to the linearized M13mp18 DNA (backbone). The small molecule antigen (FITC or Digoxigenin) was attached to one of the oligos as a 3' modification. The oligo at the 3' end of the M13 backbone was modified with a 5' biotin for attaching to streptavidin-coated micro beads in the centrifuge force microscope (CFM) experiments. The 5' region and the region complementary to the antibody-coupled oligo was left as single-stranded.

The hybridization was performed by mixing 15nM linearized M13 backbone, 150nM of each of the non-modified filler oligos, and 75nM of the biotin and antigen modified oligos, in 1x NEB Buffer 2. The mixture was set in a thermocycler with a temperature ramp from 90°C to 20°C with

a -1°C / min gradient. After the annealing, the mixture was purified using a Microspin S-400 HR column (Cytiva, 27-5140-01) to remove excess oligos. The purified linear nanoswitch construct was stored at 4°C before experiments.

Looped nanoswitch was prepared immediately before use. The antibody-coupled oligo was added to the linear nanoswitch at a 5x molar ratio in PBS buffer, with 250nM linear nanoswitch and 1250nM antibody-oligo, respectively. The mixture was incubated at room temperature for 1 hour. The looped nanoswitch was verified by gel electrophoresis.

Flow Cell Preparation

The flow cell used in the CFM experiments has single stranded DNA oligos covalently conjugated to the surface to allow DNA nanoswitches to attach via DNA hybridization. The flow cell was constructed using circular shaped cover glass (#1 thickness, 19mm diameter, Gold Seal #3346), and support glass (0.7mm thickness, 25mm diameter, S.I. Howard Glass D263), sandwiched together using double-sided Kapton tape with a flow channel shaped cut-out. Two 0.75mm diameter holes were drilled on the support glass, 15mm apart, to serve as the inlet and outlet.

Both the cover glass and the support glass were cleaned by sonication in isopropyl alcohol (ACS grade, VWR BDH1133). The glass slides were then rinsed using MilliQ water and dried in an 80°C oven.

Following the cleaning, the cover glass slides were silanized and passivated with PEG (polyethylene glycol). To prepare the cover glass slides for silanization, they were first treated by plasma etching in a SPI Plasma Prep II, under 400 milli-Torr pressure for 5 min. Next, the plasma-treated cover slides were silanized in liquid phase. 3-aminopropyltriethoxysilane (MP Biomedicals 02154766) was dissolved in acetone (Sigma 650501) to 0.5% v/v. The cover slides were placed on a holder and submerged in the silane-acetone solution, in a clean glass container. The container was then placed in a sealed desiccator and incubated at room temperature for 16 hours. Following the incubation, the cover slides were rinsed in clean acetone, then in MilliQ water, and then dried.

The silanized cover slides and cleaned support slides were then assembled into circular flow cells using Kapton double-sided tape cut to a circular outline with a flow channel shaped cut-out in the center. The flow channel was then passivated with PEG. Together with the PEG molecules, we also covalently attached ssDNA oligos to the silanized cover slides surface to serve as anchor points for the nanoswitch construct. The passivation PEG was directly attached to the silane surface via a single amine-EHS ester chemistry. The DNA oligo attachment was mediated by a dual functional linker, DBCO-PEG4-NHS, where the DBCO-PEG4-NHS was attached to the silane surface via amine-EHS ester chemistry, and then the DNA oligo with an azide active group was attached to these linkers via copper-free click chemistry. All of the reagents were mixed in a single reaction mixture and injected into the flow channel. To make

the reaction mixture, mPEG-Succinimidyl Valerate with a mean molecular weight of 5000 Da (PEG5000, Laysan Bio, MPEG-SVA-5000) was dissolved in 0.1M potassium borate buffer (Potassium Tetraborate, Sigma P5754) at pH 8.1, to a final PEG concentration of 10% w/v. The DBCO-PEG4-NHS linker (Sigma 764019) was first dissolved in DMSO (anhydrous, Invitrogen D12345) to 0.4mM, before added into the aqueous reaction mix. The anchor DNA oligo, and DBCO-PEG4-NHS, were then added into the 10% PEG5000 solution to a final concentration of 50nM DNA, and 2000nM DBCO-PEG4-NHS. 20ul of this final reaction mix solution was injected into the flow channel using a pipette. The flow cell assembly with passivation reaction mix was then placed in a petri dish sealed by Parafilm to minimize evaporation. The assembly was incubated at room temperature for 1 hour, after which the channel was flushed using 100ul MilliQ water, and incubated again with fresh passivation reaction mix for 1 hour. After the second incubation, the flow channel was flushed using MilliQ water, and dried. The flow cell assembly was kept under vacuum desiccator before use.

Immediately before each CFM experiment, the DNA nanoswitches and micro-beads were added into the flow cell channel to form tethers following the protocol in our previous work,[1] with the modification that the nanoswitches are tethered to the glass surface via DNA hybridization. Briefly, the flow cell channel was passivated by injecting 20ul of blocking solution (10 mg/ml Blocking Reagent, Roche 11096176001, in PBS), twice, with 30min incubation each at RT. The channel was then washed with 50ul experiment buffer (1x TBS, pH7.6, with 0.05% v/v Tween-20). The looped nanoswitch construct with antibody-antigen pair was diluted in the experiment buffer to 75pM. 20ul of the diluted nanoswitch solution was then injected into the channel and incubated at RT for 20min to allow hybridization to the surface. During the hybridization step, M-270 streptavidin beads (Invitrogen 65305) was washed excessively with the experiment buffer and diluted to 7mg/ml. The channel was then washed with 20ul experiment buffer, followed by injection of 20ul diluted M-270 beads to form tethers.

Electrophoresis Mobility Shift Assay

The zero-force antigen dissociation rate for the monoclonal anti-fluorescein sample was measured by electrophoresis mobility shift assay (EMSA). A fluorescein-labeled 40mer ssDNA oligo was pre-bound to unconjugated monoclonal anti-fluorescein antibody by mixing 0.1uM fluorescein-DNA to 0.5uM antibody. The mixture was incubated at RT for 10min to reach equilibrium. Following this, an excess of free fluorescein (10uM) was spiked in to compete off the fluorescein oligo and incubated at room temperature for a set amount of time (0min-60min) before loading onto a native gel. Each of the kinetic time points were timed so that all samples could be loaded onto the gel at the same time. The samples were resolved on a precast BioRad mini-protean TBE 4-20% gel, in 1x TBE buffer, at 100V, for 1hr 30min. The gel was then imaged using a Typhoon scanner in the fluorescein channel, and quantified using the gel analysis module in ImageJ. The EMSA assay was carried out in three replicates independently. The scanned gel image of one of the replicates is shown in **Fig. S1A**. The average antibody-bound

fraction of DNA was plotted against incubation time and fit to a single exponential decay function [**Fig. S1B**], which yields a dissociation lifetime T_0 = 380 ± 50 sec.

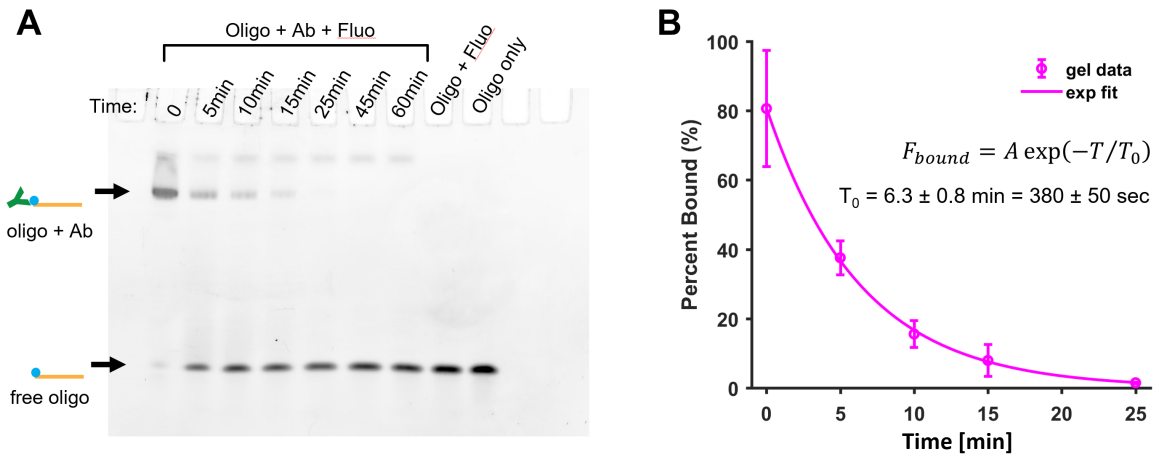

Figure S1. Electrophoresis Mobility Shift Assay of monoclonal anti-fluorescein. (A) Example gel image showing dissociation of fluorescein-oligo from monoclonal anti-fluorescein. Large excess of free fluorescein was added to the mixture to compete off the pre-bound fluorescein-oligo from the antibody. Lanes labeled as "Oligo+Ab+Fluo" shows the mixture of pre-bound fluorescein-oligo and antibody complex, plus free fluorescein spike-in, incubated at RT for the corresponding time duration. The lane labeled as "Oligo+Fluo" shows sample with free fluorescein-oligo and free fluorescein ligand. The lane labeled as "Oligo only" shows free fluorescein-oligo only. **(B)** Quantified percentage of fluorescein-oligo remained bound to the antibody (slow mobility band on the gel) versus incubation time duration from the point when free fluorescein competitor was added. The error bars indicates standard deviation from three replicate experiments. The curve is a least square fit using an exponential decay function.

Centrifuge Force Microscope Experiment

Parallel force spectroscopy experiments were carried out with a custom-built Centrifuge Force Microscope designed to fit into a benchtop centrifuge. This consisted of a miniature microscope that fit into a 3D-printed holder designed to slide into a standard centrifuge swing bucket. The table-top CFM setup and force calibration was detailed in previously published work from our group,[1] with the only modifications being a different camera (Point Grey Blackfly BFLY-PGE-50S5M-C with Sony IMX264 sensor) and a new 3D-printed swing bucket fixture to accommodate the new camera. The updated parts list is provided in Table S1. The pixel size of the new camera remains the same as in our previous work (3.4um/pixel).

Item#	Vendor	Part#	Description	Qty.
1	Thorlabs	S1LEDM	SM1-Threaded Mount for LED	1
2	Thorlabs	SM1T1	SM1 (1.035"-40) Coupler	1
3	Thorlabs	SM05RR	Retaining Ring for diffuser	1
4	Thorlabs	DG05-220	Light Source Diffuser, Ø1/2" N-BK7 Ground Glass	1
5	Thorlabs	SM1A6T	Diffuser and sample cell mount	1
6	SI Howard Glass Co	B-270	Ø 25 mm, 0.9 mm Thick	1
7	Kapton Tape	PPTDE-3	Sample Cell Assembly	1
8	VWR	63782-01	Gold Seal, #1 19 mm coverglass	1
9	Thorlabs	SM1L03	Sample Holder, SM1 Lens Tube, 0.3" Thread Depth	1
10	Thorlabs	SM1V05	Focusing Ø1" SM1 Lens Tube	1
11	Edmund Optics	#86-815	40X Olympus Plan Achromat Objective, 0.65 NA, 0.6 mm WD	1
12	Thorlabs	SM1A3	Objective Adaptor with External SM1 Threads and Internal RMS Threads	1
13	Thorlabs	AC254-100	Tube Len, f=100.0 mm, Ø1" Achromatic Doublet, ARC: 400-700 nm	1
14	Thorlabs	SM1RR	Tube Lense SM1 Retaining Ring	1
15	Thorlabs	SM1M20	Objective SM1 Lens Tube Without External Threads, 2" Long	1
16	Thorlabs	SM1A6T	Adaptor with External SM1 Threads and Internal SM05 Threads, 0.40" Thick	2
17	-	-	Custom made turning block, aluminum	1
18	Thorlabs	PFE10-P01	Turning Mirror, 1" Silver Elliptical Mirror, 450 nm - 20 μm	2
19	Thorlabs	SM1NT	Camera SM1 (1.035"-40) Locking Ring, Ø1.25" Outer Diameter	1
20	Thorlabs	SM1A9	Camera Adaptor with External C-Mount Threads and Internal SM1 Threads	1
21	Point Grey (now Teledyne FLIR)	BFLY-PGE-50S5M-C	Sony IMX264 CMOS sensor, 2448 x 2048 resolution	1
22	IMC Network	855-10734	MiniMc-Gigabit Twisted Pair to Fiber Media Converter	1
23	PrinceTel	MJX	Fiber Optic Rotary Joint	1

Table S1. Parts list for the Centrifuge Force Microscope used to carry out force spectroscopy experiments.

In this work, antibody-antigen dissociation rates were measured by constant force experiments using the CFM. After the sample flow cell was mounted in the CFM swing bucket fixture, a custom speed program was run on the centrifuge. The program started with a speed-up section at maximum acceleration, approximately 198 RPM/sec, before reaching a constant RPM speed that provided the desired constant force. The program then held the constant speed for a fixed amount of time, until most tethered nanoswitches had transitioned into the open state where the antibody-antigen pair had unbounded. The speed was then slowed down at max deceleration to a full stop. The system was held at rest for 20 min to allow the antibody-antigen pair to rebind and the nanoswitches to close again, before repeating the next force cycle. For each experiment, this constant force program was repeated for 18-20 cycles. Due to the covalent attachment chemistry we used, the bead loss over repeated force cycles was minimal [**Figure S2**]. In each repeat, a video capture of the bead motion was taken of the same region of interest in the sample flow cell [**Figure S3**A].

For the monoclonal anti-fluorescein sample, experiments were performed under 20, 25, 30, and 35 pN forces. The centrifuge RPM speed was set to 1428, 1597, 1750, and 1890 RPM, respectively. The hold time was 480, 420, 240, and 200 sec, respectively. The movies were recorded at 0.5, 0.7, 1, 1.4 fps, respectively. For both the monoclonal anti-DIG and polyclonal anti-fluorescein samples, the experiments were performed under 30pN force, with the constant centrifuge speed at 1750 RPM for 480 sec, while the movies were recorded at 0.5 fps.

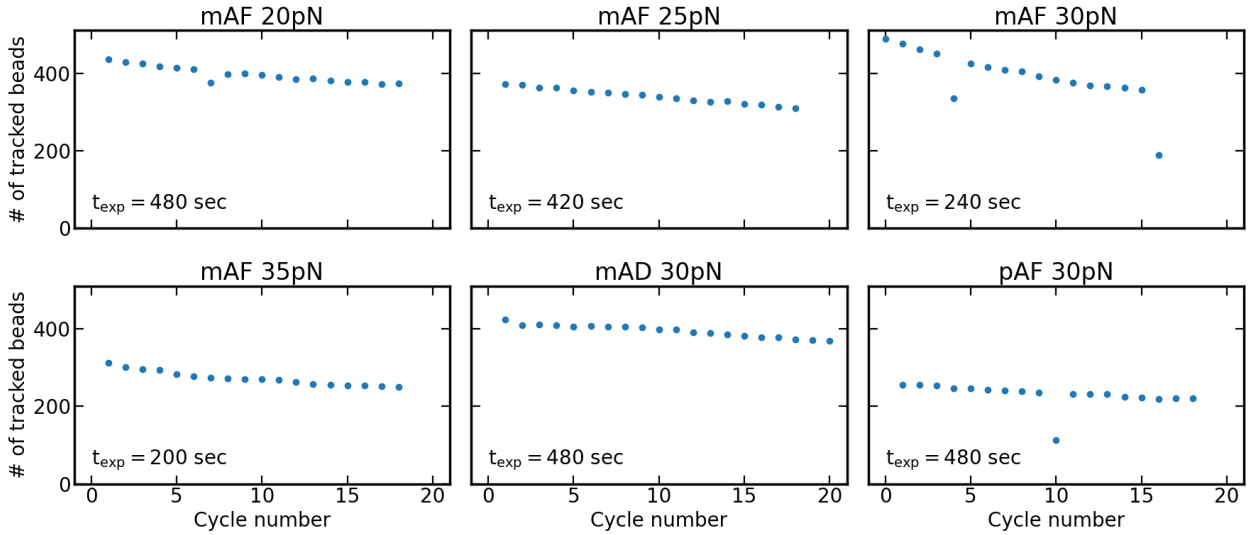

Figure S2. Number of remaining tethers per force cycle. The number of beads in the field of view of the flow cell that are successfully tracked as a function of cycle number across repeated force cycles. The hold time of the centrifuge at the indicated force, t_{exp}, is presented above for each experimental condition.

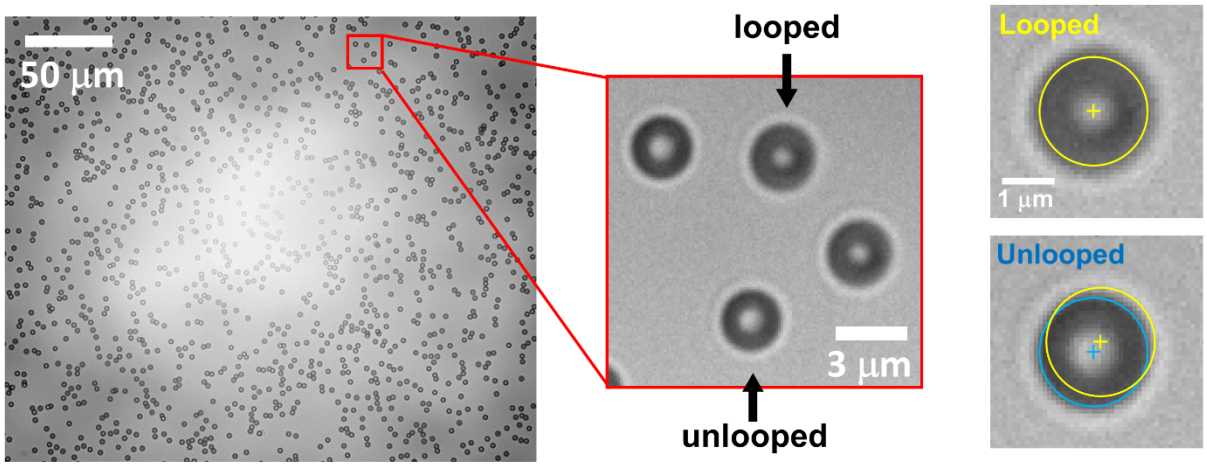

Figure S3. Raw images from the Centrifuge Force Microscope. *Left*, a portion of the camera field of view. *Center*, an inset of a few beads. The looped and unlooped nanoswitches are in slightly different focal planes. *Right*, images of the lowermost bead in the inset before and after a transition from a closed to an open nanoswitch state. The position of the center of the bead is tracked in yellow (before transition) and blue (after transition). The position of the yellow circle is overlaid in the lower right image as a reference.

Centrifuge Force Microscope Data Processing

Motion trajectories of tethered beads in drift-corrected relative x and y coordinates were extracted from the raw movie recordings as previously described.[1] Trajectories for each bead in each repeated experiment were extracted independently. The same physical beads were identified from each repeat based on the rule of minimal shift distance, and a unique Bead ID was given to each bead across repeats. In our CFM setup, the bead displacement in the y direction is the more prominent component as a result of the geometry of our CFM flow cell fixture and the swing angle of the centrifuge bucket under the working RPM speed. Thus, only the y displacement trajectories were used in the transition state analysis.

Drift-corrected y trajectories were extracted for each bead for constant force stage in each experiment cycle, with the starting point being when the spinning speed reaches a constant value. The end point was set to either the end of the stage of constant speed/force if the bead could be traced until that point, or to the point when the bead tracking was lost. A typical y trajectory shows the bead fluctuating around a stable average position under the set pulling force. A sudden transition of the mean bead position indicates that the linked antibody-antigen bound has ruptured and the nanoswitch has undergone a transition from closed to open state.

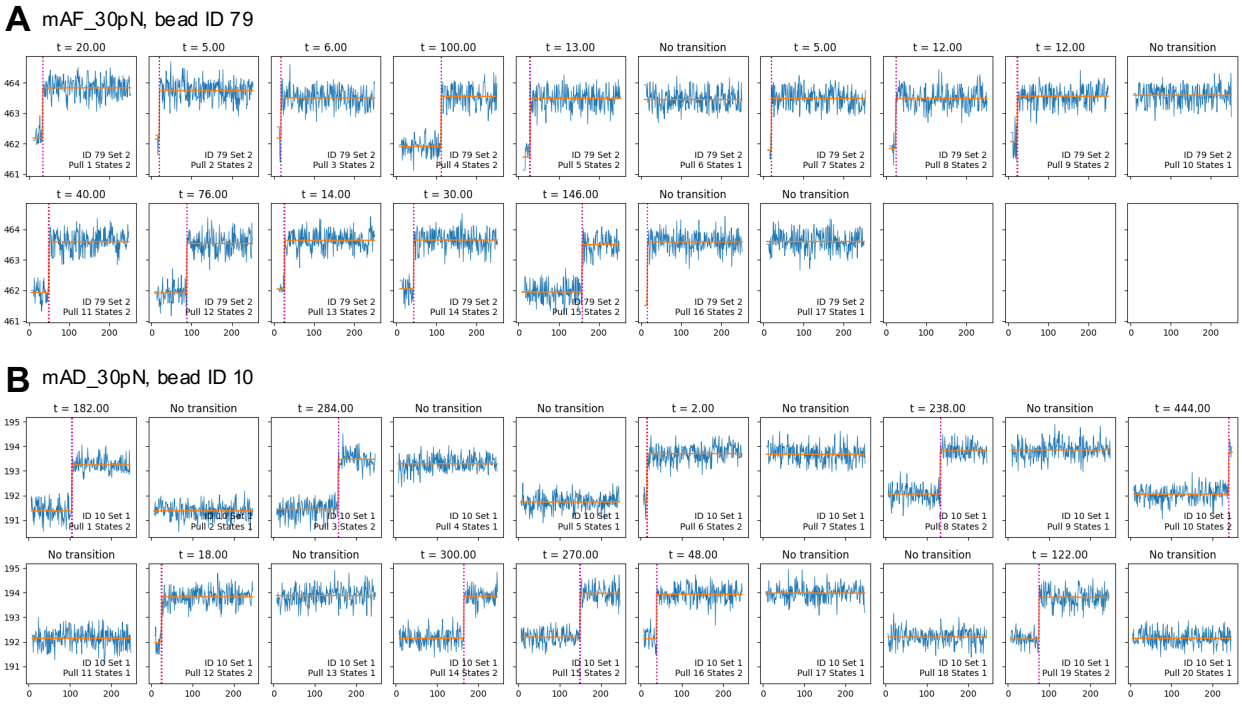

Figure S4. Example bead trajectories. An arbitrary validated bead was chosen from the **(A)** mAF_30pN dataset and from the **(B)** mAD_30pN dataset. For each of the repeated pulls where the bead was successfully tracked, the y position of the center of the bead is plotted in blue as a function of time. The trajectory fit is shown in orange. The inferred transition time is shown with the red dotted line and labeled above each plot.

These trajectories were then analyzed using the vbFRET package,[5] a Matlab package for time series analysis based on a Bayesian approach. The bead trajectories were loaded in vbFRET and fitted to idealized trajectories (step functions) using default parameters and a maximum number of states of 5. For each idealized trajectory, the number of transition events, their transition time (relative to the beginning of the trajectory) and transition distance (difference in y displacement between states) were extracted.

In addition to the transition distance and angle requirements applied to each individual transition event, as detailed in our previous work,[1] the bead transition data set was further filtered by applying of following rules on the level of the unique bead IDs, to ensure only the beads with highly reliable trajectories were included in the transition time analysis:

1. Beads from tight clusters were removed. The starting center-to-center distance matrix was obtained for each pair of beads from each repeated pull. Any bead with one or more neighbors with a bead center-to-center distance smaller than 25 pixels, or 2.2um, was removed, together with all of its neighbors within this distance range.

2. Beads with invalid transitions were removed. Only beads with zero or one upward transitions were considered valid. The upward transition is defined as the idealized y trajectory transitioning from a lower position to a higher position, corresponding to the nanoswitch going from closed to open state in our setup. Beads showing either multi-state (> 2 states) transitions, or a transition

in the incorrect direction (corresponding from open to closed state under force) were considered invalid, and all of its trajectories were removed from the data set, including repeats where it seemingly showed only valid transitions.

3. Beads with an unusually large transition distance were removed. The threshold of y transition distance dy was selected to be 3.6 pixels, or 12.2um. This threshold was determined by adding 3x standard deviation to the median dy based on all valid transitions meeting the first two requirements. Beads with one or more transition distance greater than this threshold were removed from the data set.

4. Beads with an unusually high variance in the idealized open or closed state positions between repeated pulls were removed. Beads meeting the above three requirements would show no more than two states, an open one and a closed. Consider taking the difference between the max and min values of the idealized coordinate. The larger one between the two ranges cannot be greater than a threshold of 18 pixels, or 61.2um. This threshold was set to be 5x of the dy threshold chosen in Requirement #3.

Bead IDs with one or more trajectories which failed any of the listed four requirements were removed from the data set. Overall, this additional filtering process removed ~10-30% of the remaining tracked beads. A final valid dataset of the Bead ID and transition time in each repeated measurement was compiled and used in the following Bayesian analysis.

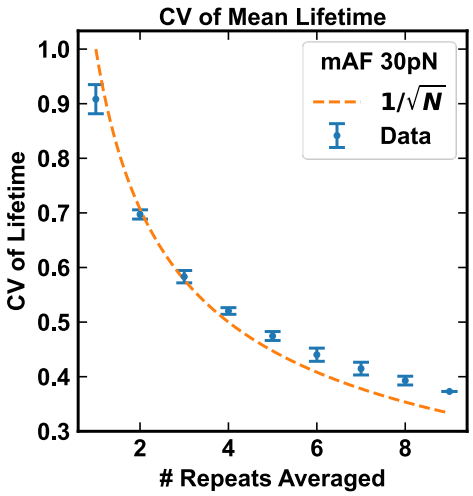

Figure S5. Dependence of lifetime uncertainty on number of repeats. For the monoclonal anti-fluorescein dataset at 30pN, all molecular pairs with at least 9 observed transitions were selected. For each molecular pair, a random subset of N transitions were randomly selected, where N is between 1 and 9, and the coefficient of variation (CV) of these transitions was calculated by dividing the standard deviation by the average. The trend line shows the $1/\sqrt{N}$ dependence expected if the only error were statistical noise.

Section S2 Bayesian Non-Parametric Fitting Method

From the table of transition times, we applied Bayesian non-parametric (BNP) inference to fit a distribution of bond-rupture lifetimes for the population of molecules. The population is modeled as a mixture of components, where the population fraction of component k is denoted p_k and the sum of the $p_k$is 1. For any molecule belonging to component k, the transition time is assumed to follow an exponential distribution with a characteristic time τ_k and a right-truncation time T of the length of the experiment, $p(t) \propto e^{-t/\tau_k}, 0 < t < T$. The BNP method determines a joint posterior distribution over the mixing proportions p_k, and the lifetimes τ_k. The posterior distribution is estimated with a Gibbs sampler Monte Carlo chain. The codes are available online at https://github.com/wpwonglab/Single-Molecule-Centrifugation.

The Weak Limit. In the non-parametric treatment, the number of components is formally infinite. Practically, however, given a finite dataset, only a small number of components have substantial posterior probability weight p_k, and the remaining components have negligible weights. We therefore apply a finite truncation and approximate the prior on $\{p_k\}$ as a Dirichlet distribution over a *finite* number of components, K_{max}. Here K_{max} should be much larger than the number of substantially populated components. In the limit $K_{max} \to \infty$, the weak limit approaches a Dirichlet Process. For our dataset and hyperparameters, we find that the posterior over $p_k$is robust for $K \geq 10$, so we are confident our truncation of the infinite Dirichlet Process is not affecting our posterior over p_k (see discussion below).

Model definition. We place an inverse gamma prior on the lifetime of each component, $\tau_k, k = 1 \dots K_{max}$, and a Dirichlet prior on the proportions p_k. Each of N beads is randomly assigned to one of the components, with probability p_k for the k'th component. The (unobserved) identity of each bead c_n is an integer between 1 and K for all beads $n = 1 \dots N$. The n'th bead has observed transitions in M_n of the repeated pulls, where $M_n \leq M$ and M is the total number of pulls. The transition time of the n'th bead on the m'th pull is t_{nm}, where $0 < t_{nm} < T,\ n = 1 \dots N, m = 1 \dots M_n$, and is assumed to be drawn from a truncated exponential distribution with a characteristic time τ_{c_n}, the lifetime of the component of that bead. To summarize,

$$\tau_k \sim InvGamma(\alpha, \beta) \quad k = 1 \dots K_{max}$$

$$\{p_1, p_2, \dots\} \sim Dirichlet(\gamma / K_{max})$$

$$c_n \sim Categorical(p_1, p_2, \dots) \quad n = 1 \dots N$$

$$t_{nm} \sim TruncExponential\left(\tau_{c_n}; T\right) \quad n = 1 \dots N, m = 1 \dots M_i$$

or explicitly,

$$P(\tau_k = \tau |\, \alpha, \beta) \propto \tau^{-(1+\alpha)} e^{-\beta/\tau}$$

$$P(p_1, p_2, \dots \mid \gamma) \propto (p_1 p_2 \dots)^{\gamma/K_{max}-1}$$

$$P(c_n = k \mid \{p_1, p_2, \dots\}) = p_k$$

$$P(t_{nm} = t \mid \{c_n, \tau_k\}) = \left[\tau\left(1 - e^{-\frac{\tau}{T}}\right)\right]^{-1} \exp\left(-t/\tau_{c_n}\right)$$

We chose an inverse gamma prior for the lifetimes because it is conjugate to the untruncated exponential distribution (which the truncated exponential distribution approaches as the cutoff time *T* increases relative to the characteristic lifetime τ_c, i.e. in the limit $T/\tau_c \to \infty$), and because it is a standard choice for biomolecular processes with first-order kinetics.[6]

An analogy for the Dirichlet Process. The key feature of the BNP formalism is that the model infers a probability distribution over $\{p_k\}$ for an infinite number of components. Intuitively, the DP can be thought of in analogy with a "Chinese restaurant", where we are tasked with seating a long line of N customers, one at a time, at a set of tables. Each customer represents a molecule and each table represents a component. As each customer enters the restaurant, they can be assigned to an existing table, with a probability proportional to the number of customers already at the table, or they can be assigned to a new singleton table, with a probability proportional to the "concentration hyperparameter" γ. As the customers are seated the number of occupied tables grows (logarithmically with N).

The hyperparameter γ determines the propensity for populating a new table. A larger value of γ induces more tables to be populated. The size of γ can be interpreted as a strength of the prior. For small γ, it is relatively rare to populate a new table, so it only tends to happen when a bead is highly distinct from existing data points such that it populates its own table. By contrast, for large γ, it is relatively more likely to populate new tables, and so tables are populated even when a customer can comfortably be assigned to one of the existing tables. Here the preference of a customer for a pre-populated table represents the likelihood of observing a bead's transition times given the lifetime of an existing component (as compared to forming a new component with a distinct lifetime).

Prior and Posterior. The parameters we wish to determine are the mixing proportions p_k, and the lifetimes τ_k. As with any Bayesian method, we assume a prior distribution over these parameters, encoded in the first two equations above. These represent our existing knowledge about these parameters.

The data is used to infer a posterior distribution over these parameters, which can be interpreted as the model's confidence that the parameters have a set of values based on the dataset. From the maximum and spread of this posterior distribution over p_k and τ_k, we can infer point estimates and relative errors of the parameters. In our case, the posterior is numerically sampled with a Monte Carlo approach.

Gibbs Monte Carlo sampler. In the first iteration, the parameter values are drawn from the prior distribution. Afterwards, the value of the parameters at each iteration is determined stochastically from the values at the previous iteration, as described shortly. Each iteration is a

sample. After many samples, the distribution of parameter values approaches the posterior. In each iteration t, we estimate a set of putative parameter values, $p_k^{(t)}$ and $\tau_k^{(t)}, k = 1 \ldots K$, as well as a putative way to assign beads to component, $c_n^{(t)}, n = 1 \ldots N$, based on the values of variables in the previous iteration , $p_k^{(t-1)}, \tau_k^{(t-1)}$, and $c_n^{(t-1)}$. A schematic explanation of the Gibbs sampler is given here, and a more complete description is given at the end of this section.

1. Update the component assigned to each bead $c_n^{(t)}, i = 1 \ldots N$, based on the values of $p_k^{(t-1)}$ and $\tau_k^{(t-1)}$ from the previous iteration.
2. Update the proportion of beads in each component, $p_k^{(t)}, k = 1 \ldots K$, based on the number of beads currently assigned to that component, determined from $c_n^{(t)}$.
3. Update the lifetime of each component, $\tau_k^{(t)}, k = 1 \ldots K$, based on the observed transition times of all beads currently assigned to that component, determined from $c_n^{(t)}$.

The new set of values for the variables $p_k^{(t)}, \tau_k^{(t)}$, and $c_n^{(t)}$ are saved to the computer, and the procedure is repeated.

Sampling of Dirichlet Process. In the weak limit, the mixing proportions p_k are sampled as a standard Dirichlet distribution over K_{max} classes. Even though p_k is nonzero for all k, the probability weight ends up getting concentrated in a small number of components. Therefore, when the bead identities c_n are sampled (with standard multinomial sampling), they end up being assigned to a few components, and the remainder of the components are empty. As the Monte Carlo chain proceeds, the number of beads assigned to each component goes up and down, as shown in **Figure S6**. Sometimes a component loses all of its beads; conversely, sometimes an empty component begins to accumulate beads. (At any iteration, all empty components are assigned a random lifetime drawn from the prior distribution.)

To get a sense of the number of significant subpopulations in our biomolecular sample, we count the number of components with at least one bead. This is something which is (indirectly) sampled over the course of the Monte Carlo chain. We ran the sampling for long enough to attain a steady-state distribution of the number of significant subpopulations. (As a technical note, it is difficult to directly interpret this as the number of distinct kinetic subpopulations, because by the nature of the Dirichlet Process, the number of populated mixture components is expected to grow logarithmically with the number of data points.)

Since component lifetimes τ_k have a non-conjugate prior, there is no standard method for sampling from the posterior. We used a rejection sampling approach: the graph of the conditional posterior density $p(\tau_k|c_n)$ is plotted as a function of τ_k. A point (x, y) is chosen uniformly and at random with $1 < x < 500$ and $0 < y < \max p(\tau_k|c_n)$. If the point lies under the graph of the probability density (i.e., $y < p(x)$), then x is returned as a sample of τ_k; if not, the procedure is repeated. The first value of x which lies under the graph is returned.

Hyperparameters of the prior. The results of the Bayesian fit can be sensitive to the hyperparameters α and β for the lifetime prior and γ for the DP. We chose biophysically plausible values of $\alpha = 2, \beta = 240$, which cover the range of lifetimes exhibited by our monoclonal antibody samples. The optimal value of the concentration hyperparameter γ was determined by testing different values and inspecting the results, as shown in **Figure S7**. As expected, the larger values of γ induced a larger posterior probability for populating more mixture components, whereas smaller values of γ had smaller number of populated mixture components. We ended up choosing a value of $\gamma = 10^{-3}$; the range from $\gamma = 3 \times 10^{-4} - 3 \times 10^{-3}$ led to similar results (see Figure S5).

Sampling scheme. We sampled 60,000 iterations, discarded the first 10,000, and took every 100th sample after that. As shown in **Figure S8B**, the posterior has converged after 1,000 samples, and the correlation between samples drops off after 100 samples.

The Monte Carlo chain should not be sensitive to the initialization as long as the samples before equilibration are discarded. We tested this by initializing in two ways, (1) with all beads assigned to the same component, and (2) with equal probability in all K_{max} component. Both initial conditions led to both the same equilibrium after several hundred iterations, as shown in **Figure S6**.

We also verified that our finite truncation is justified. **Figure S8** shows $K_{max} = 10$ and **Figure S9** shows $K_{max} = 50$. The results are nearly identical.

Full algorithm. The full algorithm is described below. In each step, the posterior is written in the form (prior) × (likelihood).

0. Initialize the values of p_k, τ_k, and c_n.

Repeat until convergence:

1. For each bead $n = 1 \dots N$, update bead assignment c_i by sampling from

$$P(c_n = k \mid \{p_k, \tau_k, t_{nm}\}) = \frac{L_k}{\sum_{l=1}^{K_{max}} L_l}, \quad \text{where}$$

$$L_k = p_k \times \prod_{m=1}^{M_n} \left[\tau_k \left(1 - e^{-\frac{\tau_k}{T}}\right)\right]^{-1} \exp\left(-t_{nm}/\tau_k\right) \quad \text{for } k = 1 \dots K_{max}$$

2. For each component, $k = 1 \dots K$, update the mixing proportions:

$$P\left(p_1, \dots, p_{K_{max}} \mid \{c_i, \tau_c, t_{ij}\}\right) \propto (p_1 p_2 \cdots p_K)^{\gamma - 1} \times \prod_{k=1}^{K} p_k^{n_c}, \quad \text{where}$$

$$n_c = \sum_{i=1}^{N} \delta(c_i, c)$$

3. For each component $k = 1 \ldots K$, update the lifetime τ_k based on the observed transition times of all beads currently classified as that component. This is done with a rejection sampling approach. For unoccupied components, the lifetime is drawn from the prior.

$$P(\tau_k = \tau \mid \{p_k, c_n, t_{nm}\}) \propto \left(\tau^{-(1+\alpha)} e^{-\beta/\tau}\right) \times \prod_{n=1}^{N} \prod_{m=1}^{M} \delta(c_n, c) \left[\tau \left(1 - e^{-\frac{t_{nm}}{T}}\right)\right]^{-1} e^{-t_{nm}/\tau}.$$

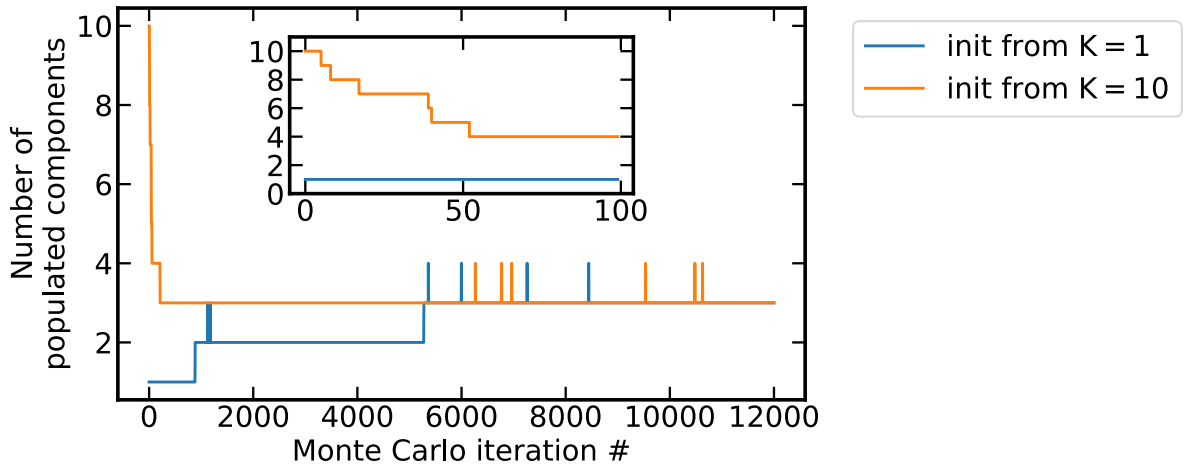

Figure S6. Comparison of different initialization schemes for BNP. The number of non-empty components is plotted as a function of the Monte Carlo iteration for an initialization with K=1 component (blue) and with K=10 components (orange). Inset, a zoom-in for early Monte Carlo iterations.

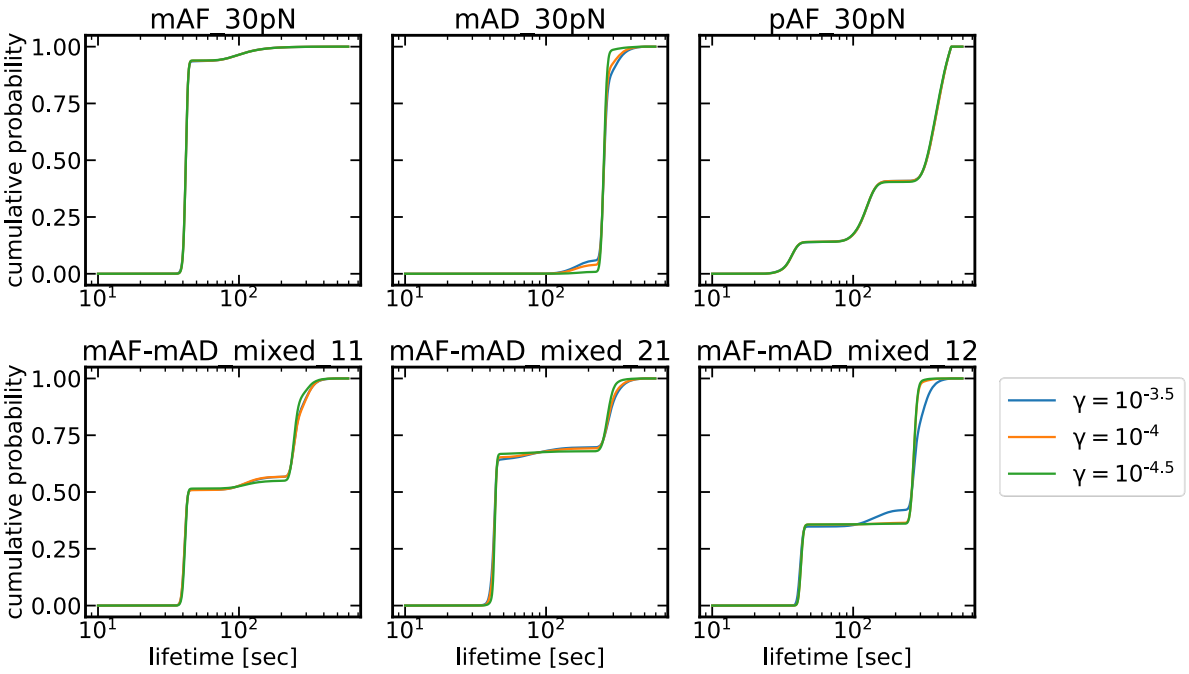

Figure S7. Comparison of different values of the concentration hyperparameter, gamma. Plotted is the cumulative posterior probability density of the unbinding lifetime, inferred with various values of the concentration hyperparameter.

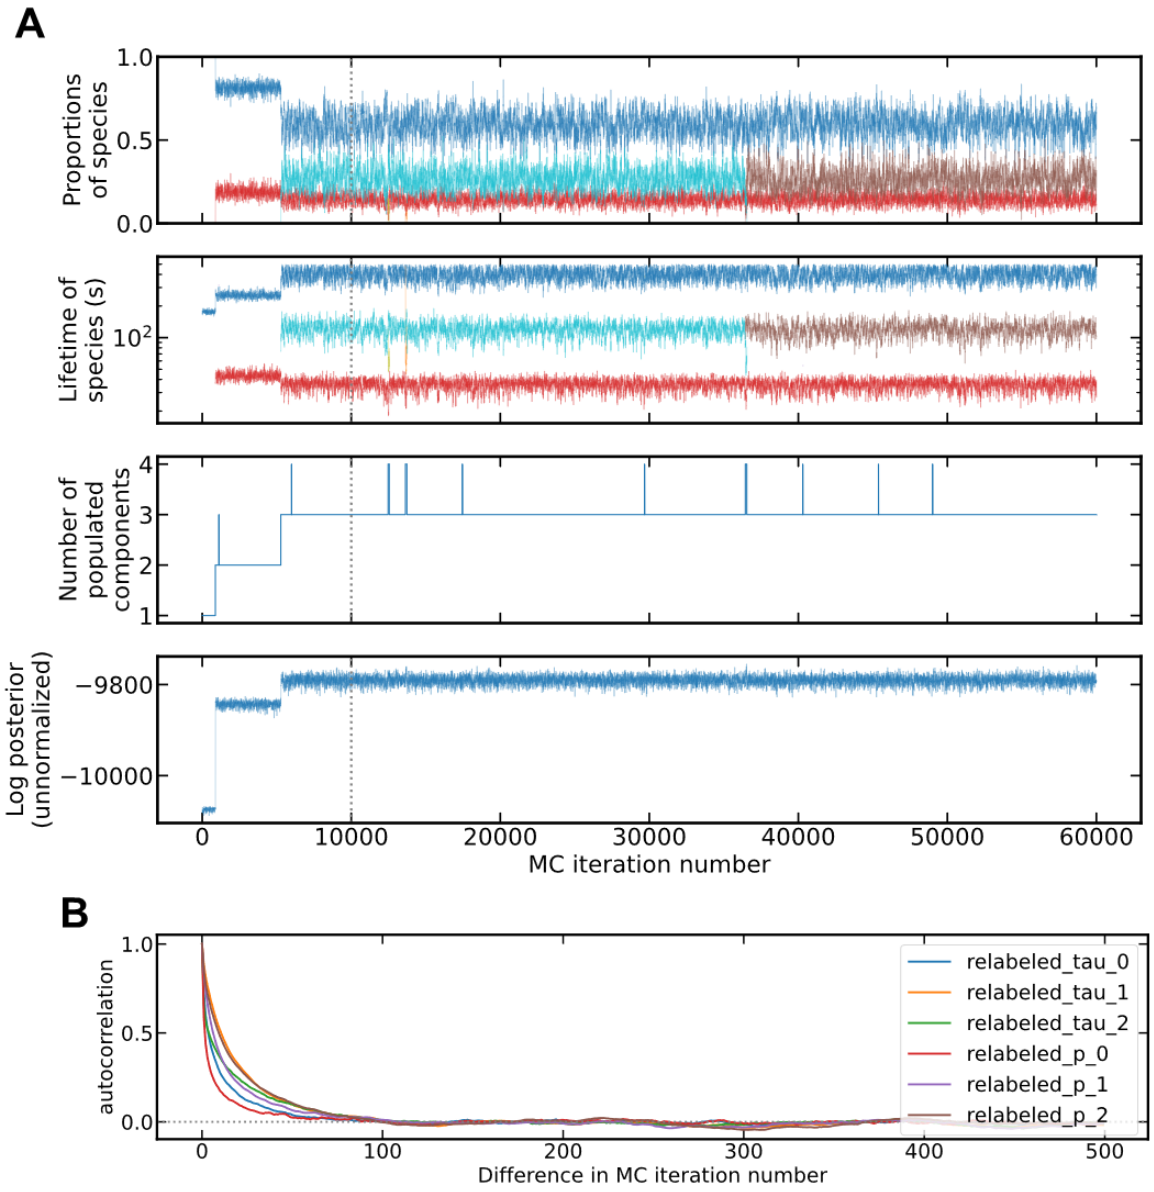

Figure S8. Convergence of Monte Carlo chain for the pAF dataset, with $K_{max} = 10$. (A) The values of various parameters are tracked as a function of Monte Carlo iteration number. The top two panels track p_k and τ_k, respectively, with each component labeled with a different color. Dotted vertical line shows 10,000 iterations, the equilibration time. **(B)** The autocorrelation of parameters, after relabeling the k from lowest to highest lifetime to resolve label-switching ambiguities. The first 10,000 iterations are not included in the calculation of autocorrelation.

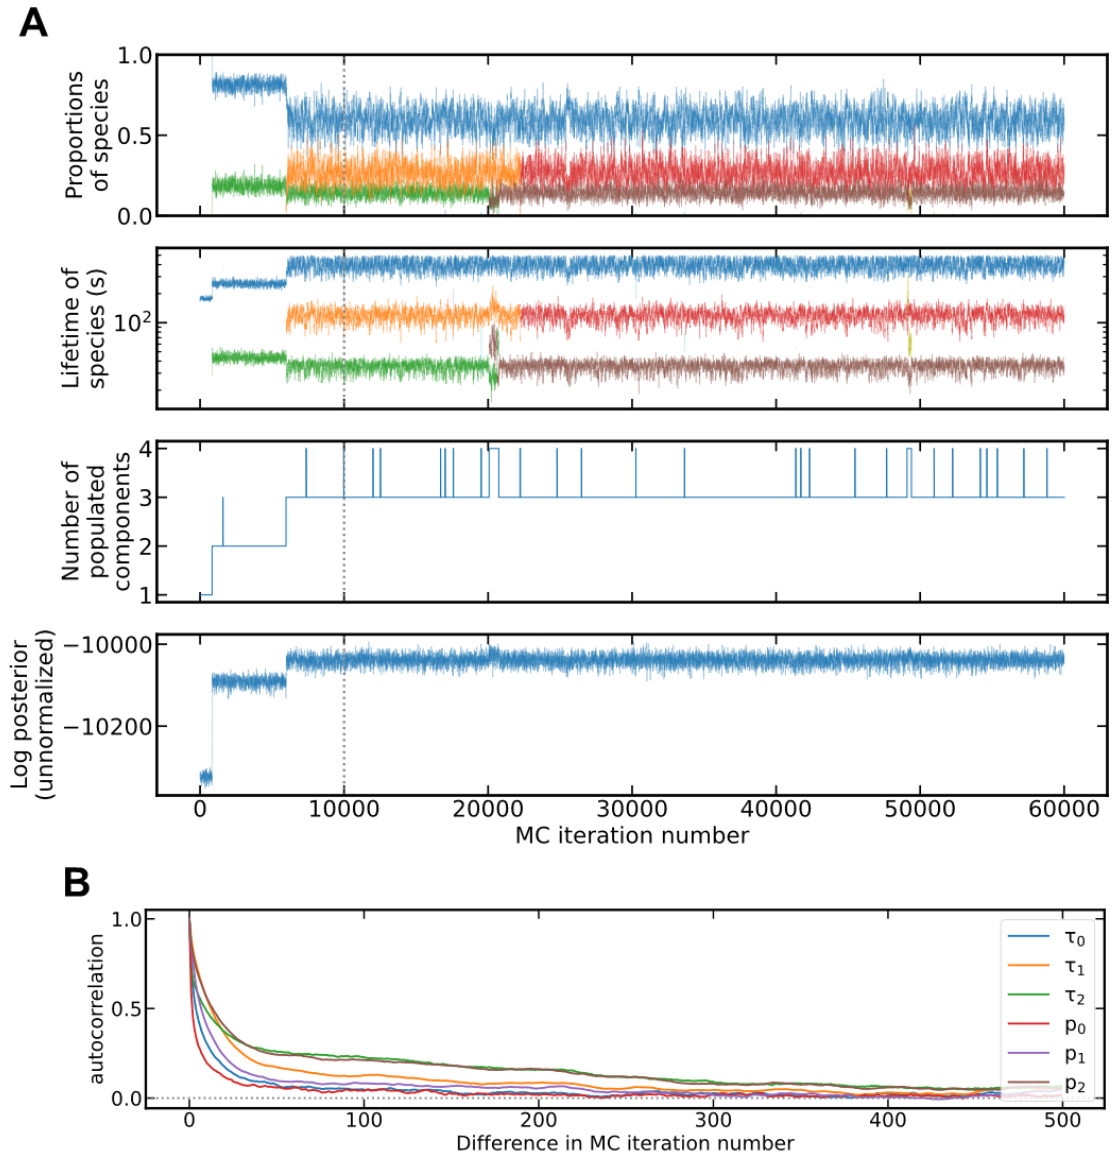

Figure S9. Convergence of Monte Carlo chain for the pAF dataset, with $K_{max} = 50$. Identical to Figure S4 except with K_max=50. Note that the results are nearly identical to Figure S4, indicating that the finite truncation of the Dirichlet Process is acceptable.

Section S3 How per-molecule information can help resolve components

The fitting of multi-exponential decay curves is a common task in science and engineering. In most situations, each data point is independent of the others and there is no information about whether any two given data points originate from the same mixture component. However, since our experimental method tracks the same molecule between repeated CFM pulls, we have access to an extra piece of information: namely, all the data points coming from the same molecule must be in the same component.

To investigate how this per-molecule information could aid in parameter fitting, we tested our analysis method with a standard parameter set, BOLIDEN3, which has been used as a benchmark for assessing fitting methods for multi-exponentials.[7–9] Prior work has shown that the Bayesian non-parametric approach can successfully infer BOLIDEN3 parameters with as few as 10^4 data points[9].

To assess the utility of per-molecule information, we generated two datasets:

1. 10,000 molecules with 1 pull each with BOLIDEN3 parameters (**Figure S10A**)
2. 500 molecules with 20 pulls each with BOLIDEN3 parameters (**Figure S10B**)

Both datasets have the same total number of data points, comparable to what is feasible with the CFM. The difference is that dataset 2 has additional per-molecule information.

Shown in **Figure S10B** is the per-molecule averaged histogram of dataset 2, analogous to the inset in **Figure 2C** of the main text. Visually, the per-molecule averaging in dataset 2 allows peaks to become clearly resolved. Applying our Bayesian nonparametric analysis to these datasets (hyperparameters $\alpha = 0.1, \beta = 0.5$) yields the parameter estimates in **Table S2.** Interestingly, dataset 2 exhibits a smaller uncertainty in the lifetime estimate than dataset 1, owing to the extra information (purple vs blue text). However, the uncertainty in mixing proportion in dataset 2 is larger because of the fewer total number of molecules. Hence, for a given total number of data points, the repeated per-molecule measurement can indeed reduce the statistical uncertainty in the lifetime estimate (at the cost of greater uncertainty in the proportion estimate).

		tau_1	tau_2	tau_3	tau_4	p_1	p_2	p_3	p_4
parameter value		142.83	9.46	1.362	0.1000	0.750	0.174	0.0587	0.0167
dataset 1	mean	140.85	9.34	1.439	0.0795	0.750	0.172	0.0625	0.0141
	std	2.44	1.62	0.489	0.0337	0.017	0.025	0.0215	0.0059
dataset 2	mean	142.06	9.82	1.454	0.0953	0.747	0.177	0.0580	0.0160
	std	1.64	0.23	0.060	0.0075	0.060	0.053	0.0326	0.0176

Table S2. Results of applying Bayesian Nonparametric analysis to synthetic datasets with the BOLIDEN 3 parameters.[7] The uncertainty in lifetime for dataset 1 and dataset 2 are shown in purple and blue, respectively.

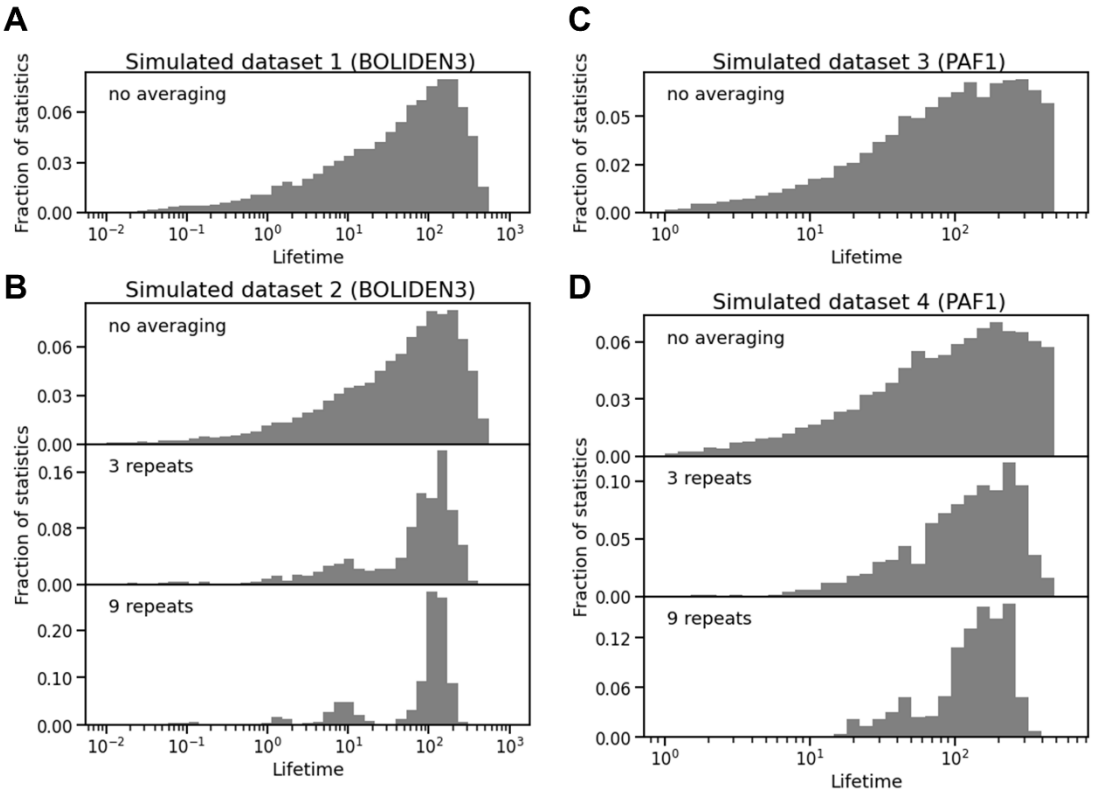

Figure S10. Effect of per-molecule averaging for simulated datasets. (A) Simulated dataset 1 with BOLIDEN 3 parameters with 10,000 molecules and 1 repeat each. **(B)** Dataset 2 with BOLIDEN3 parameters with 500 molecules and 20 repeats each. Top, middle, and lower panels show with averaging over differing numbers of molecules. **(C)** Simulated dataset 3; same as (A) but with the PAF1 parameters. **(D)** Simulated dataset 4; same as (B) but with PAF1 parameters. Note the similarity to **Fig. 4B** in the main text.

To see whether this result also holds for a more biophysically plausible parameter set, we repeated the above analysis using the parameters we inferred from the polyclonal anti-fluorescein sample. These parameters, which we call PAF1, are listed in Table 2 in the main text. As before, we simulated two datasets:

3. 10,000 molecules with 1 pull each with PAF1 parameters (**Figure S10C**)
4. 500 molecules with 20 pulls each with PAF1 parameters (**Figure S10D**)

To emulate the experimental conditions, we removed data points longer than the truncation time of 480 seconds.

As shown in **Figure S10D**, the per-molecule averaged histograms of the PAF1 simulated dataset 4 bear a close resemblance to the histograms of the experimental data (**Fig. 4B**), indicating that our fitted parameters largely recapitulate the original data. While our analysis suggests that three components may be sufficient to explain the observed unbinding kinetics, we note that due to the ill-conditioned nature of the multi-exponential, the data could also be well-modeled by a greater number of components or even by a continuous lifetime density. In

this case, the parameter estimation is greatly aided by the repeated measurement: without this information, dataset 3 does not converge to the correct values, whereas with this added information, dataset 4 converges to the true parameter values (**Table S3**).

		tau_1	tau_2	tau_3	p_1	p_2	p_3
parameter value		400.079	119.566	35.451	0.594	0.266	0.140
dataset 3	mean	300.332	156.485	46.212	0.589	0.155	0.254
	std	58.636	75.534	8.910	0.263	0.271	0.074
dataset 4	mean	403.510	122.435	36.489	0.566	0.287	0.145
	std	31.977	5.849	1.092	0.076	0.071	0.049

Table S3. Results of applying Bayesian Nonparametric analysis to synthetic datasets with the PAF1 parameters. The uncertainty in lifetime for dataset 3 and dataset 4 are shown in purple and blue, respectively.

Section S4 References

(1) Yang, D.; Ward, A.; Halvorsen, K.; Wong, W. P. Multiplexed Single-Molecule Force Spectroscopy Using a Centrifuge. *Nat. Commun.* **2016**, *7* (1), 11026. https://doi.org/10.1038/ncomms11026.

(2) Yang, D.; Wong, W. P. Repurposing a Benchtop Centrifuge for High-Throughput Single-Molecule Force Spectroscopy. In *Single Molecule Analysis*; Peterman, E. J. G., Ed.; Methods in Molecular Biology; Springer New York: New York, NY, 2018; Vol. 1665, pp 353–366. https://doi.org/10.1007/978-1-4939-7271-5_19.

(3) Halvorsen, K.; Schaak, D.; Wong, W. P. Nanoengineering a Single-Molecule Mechanical Switch Using DNA Self-Assembly. *Nanotechnology* **2011**, *22* (49), 494005. https://doi.org/10.1088/0957-4484/22/49/494005.

(4) Koussa, M. A.; Halvorsen, K.; Ward, A.; Wong, W. P. DNA Nanoswitches: A Quantitative Platform for Gel-Based Biomolecular Interaction Analysis. *Nat. Methods* **2015**, *12* (2), 123–126. https://doi.org/10.1038/nmeth.3209.

(5) Bronson, J. E.; Fei, J.; Hofman, J. M.; Gonzalez, R. L.; Wiggins, C. H. Learning Rates and States from Biophysical Time Series: A Bayesian Approach to Model Selection and Single-Molecule FRET Data. *Biophys. J.* **2009**, *97* (12), 3196–3205. https://doi.org/10.1016/j.bpj.2009.09.031.

(6) English, B. P.; Min, W.; van Oijen, A. M.; Lee, K. T.; Luo, G.; Sun, H.; Cherayil, B. J.; Kou, S. C.; Xie, X. S. Ever-Fluctuating Single Enzyme Molecules: Michaelis-Menten Equation Revisited. *Nat. Chem. Biol.* **2006**, *2* (2), 87–94. https://doi.org/10.1038/nchembio759.

(7) Watson, G. A. Chebyshev Approsimation to Data by Positive Sums of Exponentials. *IMA J. Numer. Anal.* **1990**, *10*, 569–582.

(8) Landowne, D.; Yuan, B.; Magleby, K. L. Exponential Sum-Fitting of Dwell-Time Distributions without Specifying Starting Parameters. *Biophys. J.* **2013**, *104* (11), 2383–2391. https://doi.org/10.1016/j.bpj.2013.04.030.

(9) Hines, K. E.; Bankston, J. R.; Aldrich, R. W. Analyzing Single-Molecule Time Series via Nonparametric Bayesian Inference. *Biophys. J.* **2015**, *108* (3), 540–556. https://doi.org/10.1016/j.bpj.2014.12.016.