A general overview of

Free energy methods

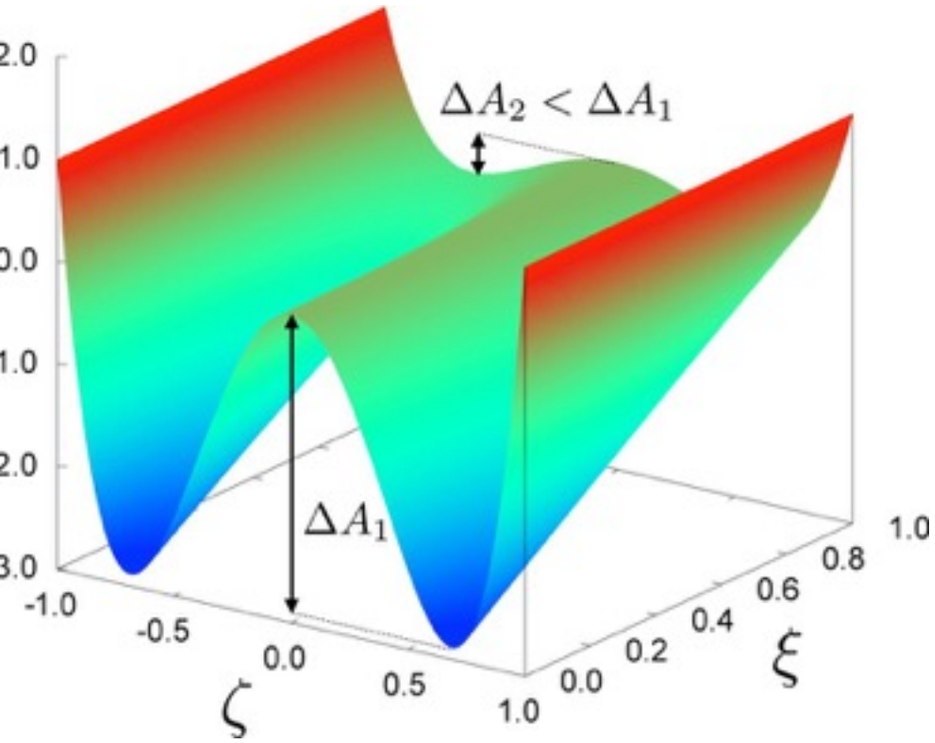

Comer et al. (2014)
JPCB 119:1129.

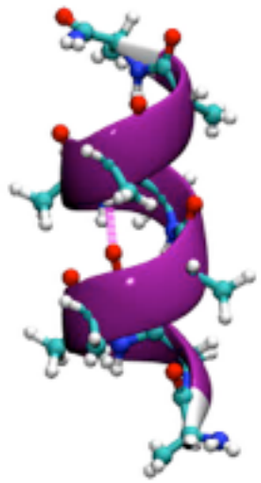

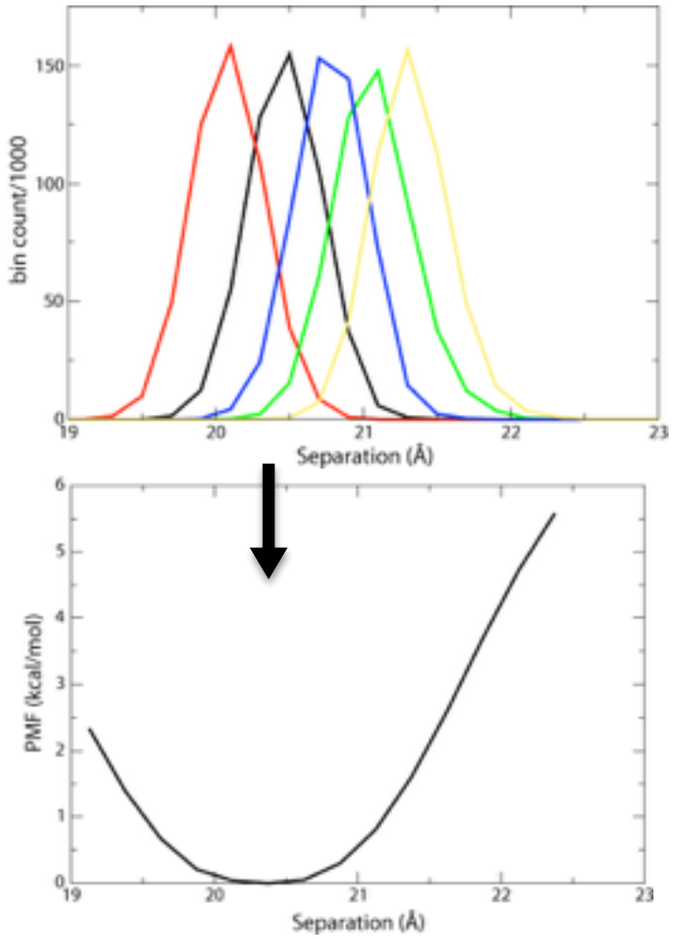

James C. (JC) Gumbart
Georgia Institute of Technology, Atlanta

Computational Biophysics Workshop | Georgia Tech | Nov. 16 2016

Outline

I. *Introduction to free energy and the need for biases*

II. *Classes of common free-energy calculation methods*

III. *Choosing a reaction coordinate and what is left behind*

IV. *Example 1: deca-alanine folding*

V. *Example 2: membrane permeation*

VI. *Final thoughts*

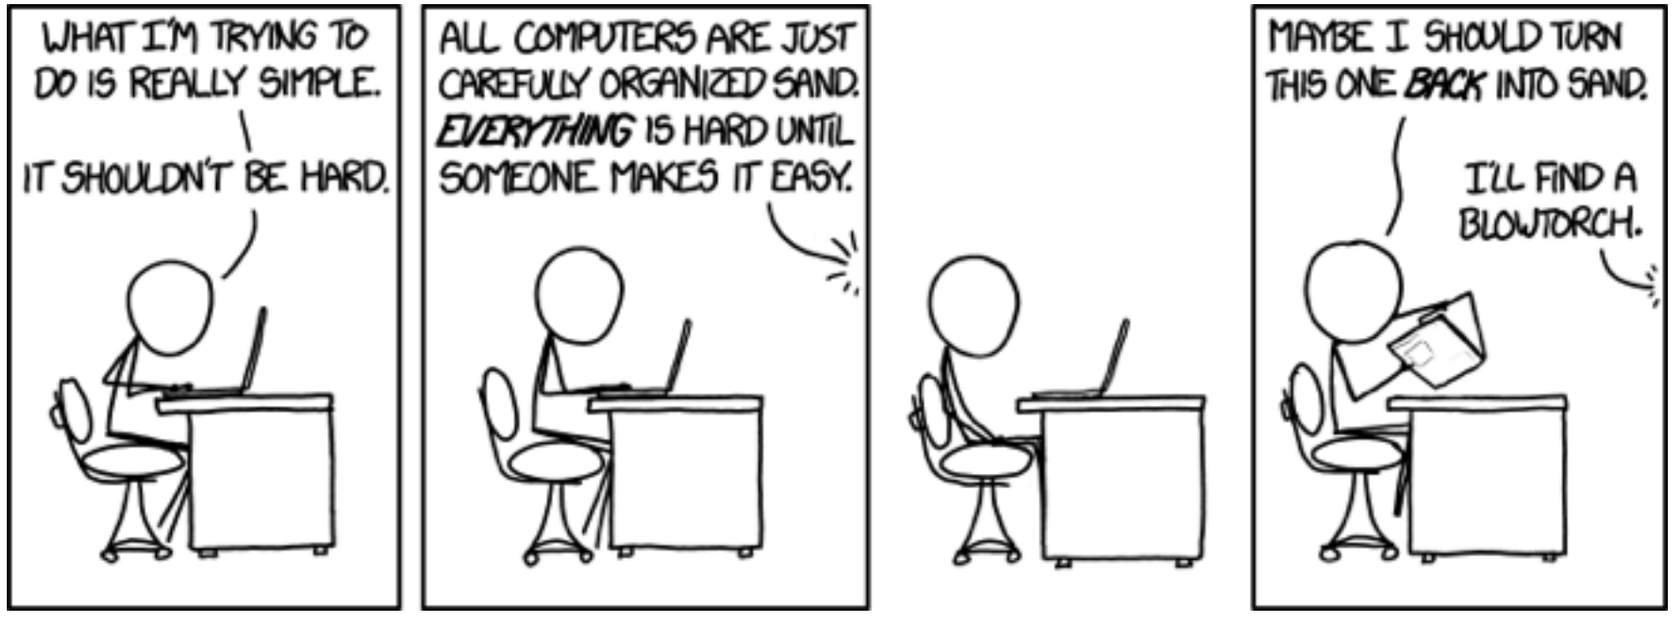

http://xkcd.com/

What is “free energy”?

the energy available to do work

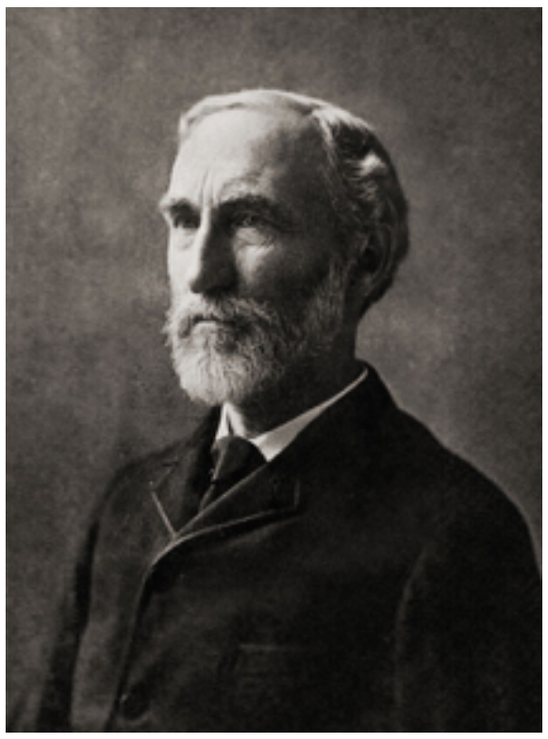

Josiah Gibbs (technically, a mathematician)

If you’re a physicist…

$$\Delta A = \Delta U - T\Delta S$$

***common** view of “energy”*

*entropy - an **effective** energy*

constant volume/gas phase

If you’re a chemist…

$$\Delta G = \Delta U - T\Delta S + p\Delta V$$

constant pressure

Hermann von Helmholtz

International Union of Pure and Applied Chemistry (IUPAC) has been trying to eliminate the term for years!

Fugacity: Fugacity is a thermodynamic quantity related to the Gibbs energy (older terminology of Gibbs free energy not recommended) of a gaseous molecule. For pressures of gases near 1 atm, the partial pressure of

from IUPAC Glossary of atmospheric chemistry terms 1990.

What can we measure with free energy?

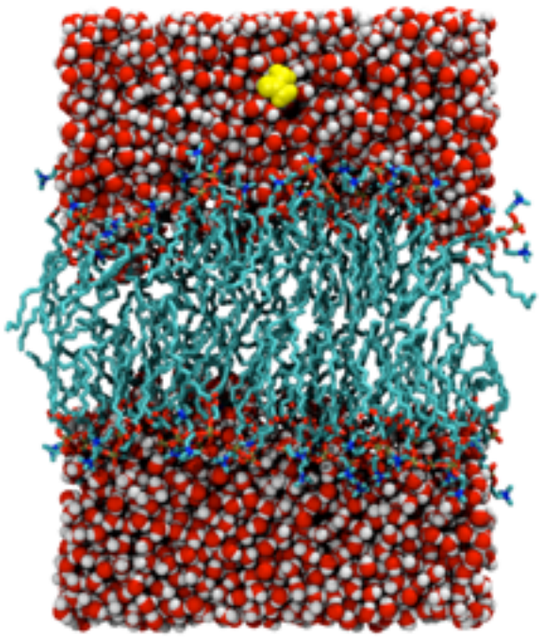

partition coefficients

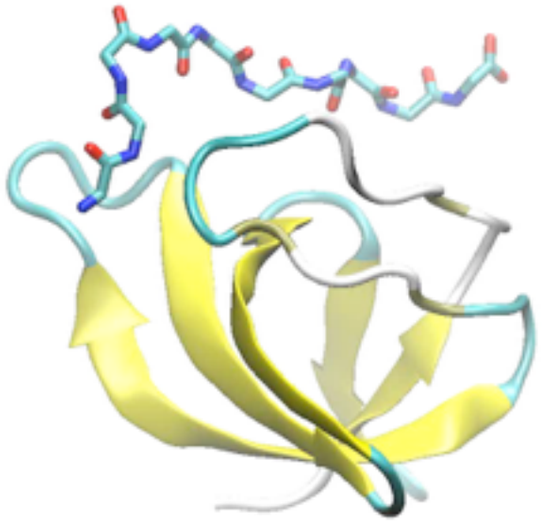

binding energies

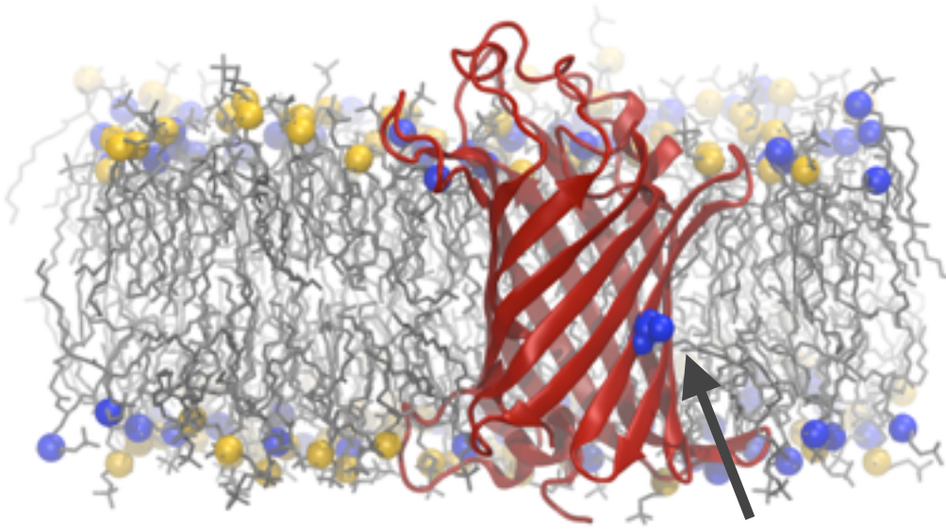

mutagenesis effects

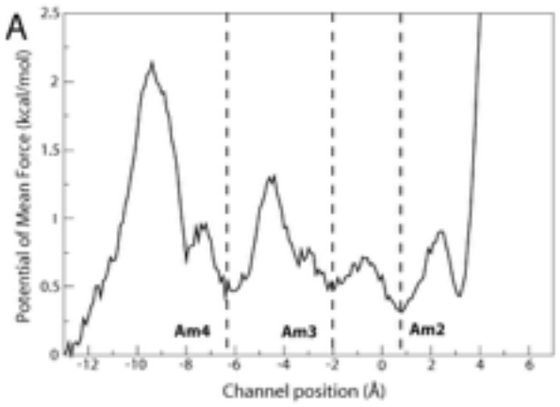

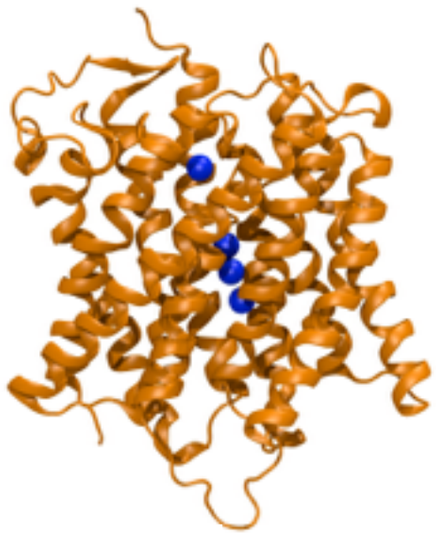

activation barriers

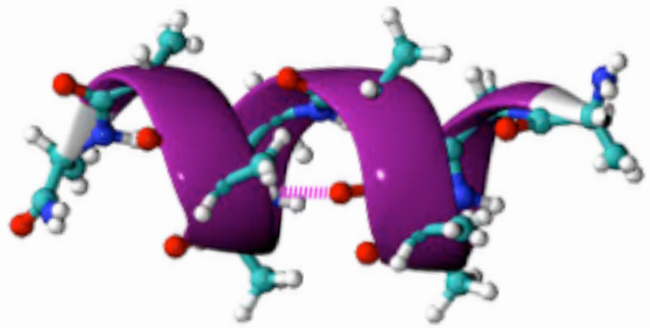

conformational change

How to measure free energy?

$$\beta = \frac{1}{kT}$$

Free energy tells us something about a system that is **general**, rather than specific to a single experiment or simulation

$$\Delta A = -\frac{1}{\beta}\ln(\frac{Z_1}{Z_0})$$ ← partition functions for two different states of the system

$$p_i \propto Z_i \longrightarrow \Delta A = -\frac{1}{\beta}\ln(\frac{p_1}{p_2})$$ probability of observing a given state

$$\lim_{t\to\infty}\frac{1}{t}\int q(t)dt = \frac{1}{M}\sum_M q_i = \langle q \rangle$$ Ergodic hypothesis - running a simulation is the same as doing an experiment many times

Therefore, we can determine free energies just by running the experiment/simulation for a sufficiently long time and counting how often a given state appears!

The value of brute-force simulations…

10^8

10^6

10^4

10^3

atoms

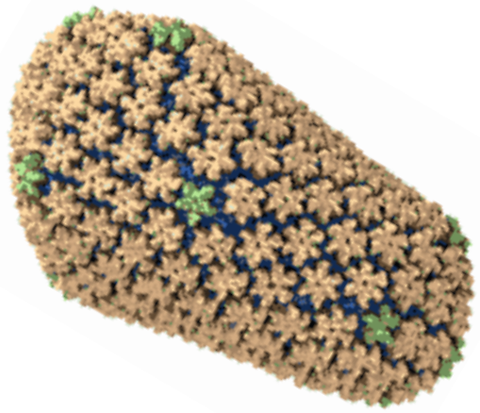

HIV capsid:
200 ns
(2013)

Zhao et al. Mature HIV-1 capsid structure by cryo-electron microscopy and all-atom molecular dynamics. (2013) *Nature* 497:643-6.

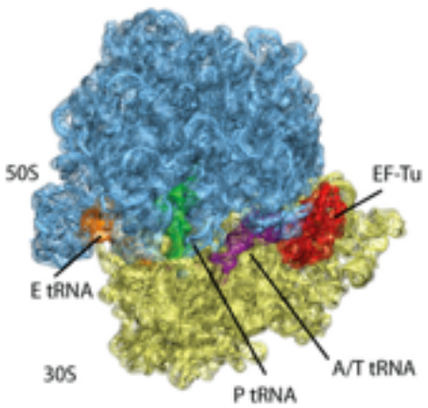

ribosome:
20 ns
(2005)

Sanbonmatsu, Joseph, Tung. Simulating movement of tRNA into the ribosome during decoding. (2005) *PNAS* 102:15854–59.

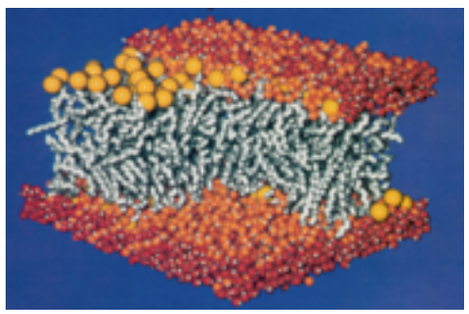

membrane:
363 ps
(1993)

Heller, Schaefer, Schulten. Molecular dynamics simulation of a bilayer of 200 lipids in the gel and in the liquid crystal-phases. (1993) *J. Phys. Chem.* 97:8343-60.

BPTI:
8 ps
(1977)

McCammon, Gelin, Karplus. Dynamics of folded proteins. (1977) *Nature* 267:585-90.

…is not time!

What is the lifetime of a graduate student?

Or, in other words, how long would you like to wait?

(isoproterenol, used to treat bradycardia and heart block).
We also simulated the binding of dihydroalprenolol to the β_1-adrenergic receptor (β_1AR). We performed 82 simulations lasting from 1 to 19 μs each, resulting in a total of 21 spontaneous binding events (Tables S1 and S2).

Using **Anton**, over 200 μs of simulation were needed to observe multiple binding events

On a normal supercomputer at even 100 ns/day, would need 5.5 years! (So one grad student per ligand…)

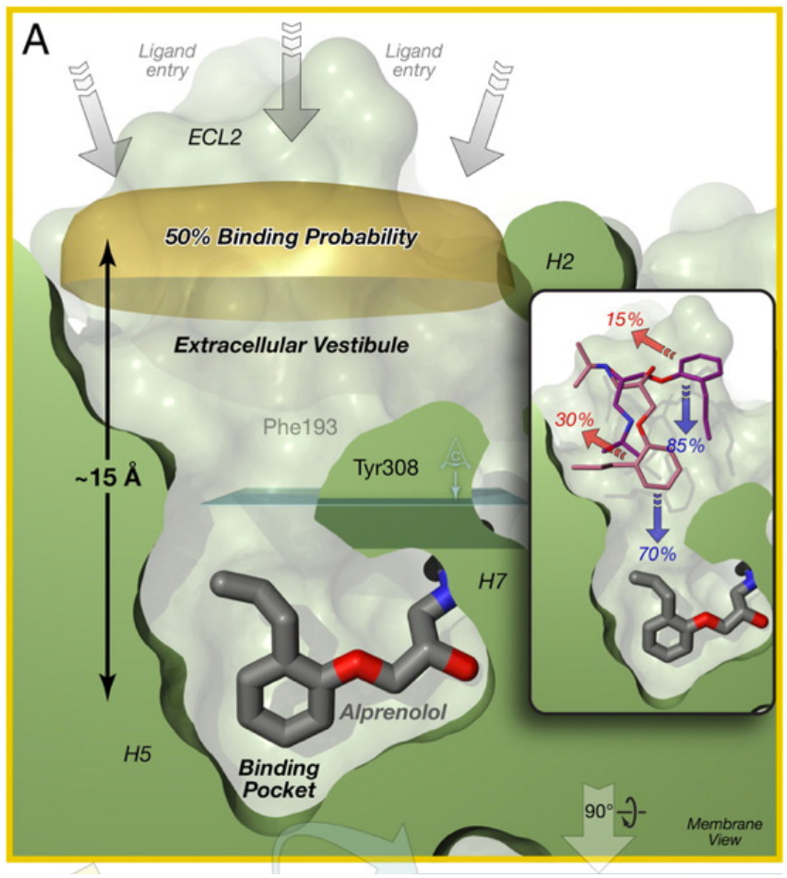

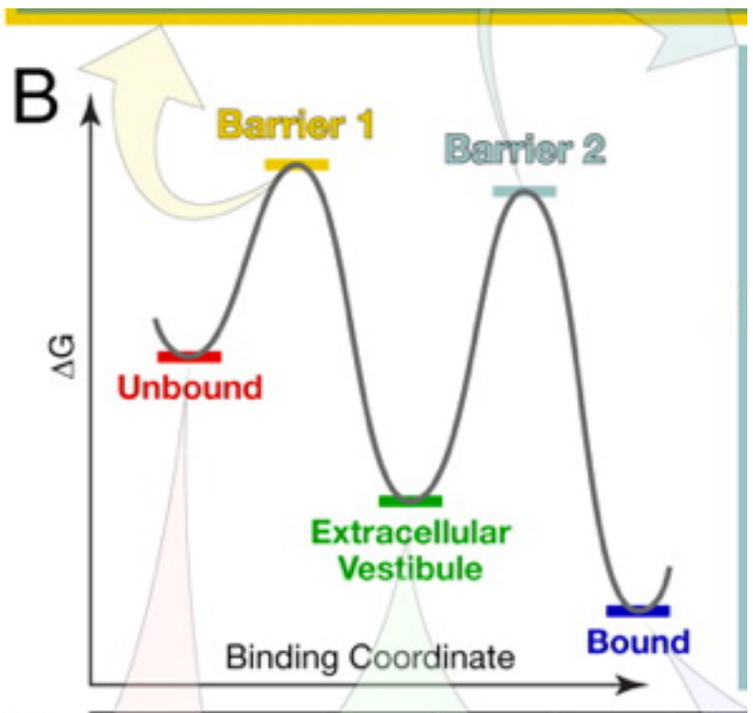

BUT!!! Notice that there are no units?

To get free energy, must observe **UNbinding** as well as **binding**, a process orders of magnitude slower!

Dror,…,Shaw. Pathway and mechanism of drug binding to G-protein-coupled receptors. (2011) *PNAS* 108:13118-123.

How can we speed up the process?

Why don't we just ***heat**** *up the system?*

≠

We are not sampling the correct thermodynamic ensemble if we just unnaturally perturb the system!

We must **enhance** the sampling along a particular coordinate or coordinates (*can be collective*) in such a way that we can **recover** the correct thermodynamics

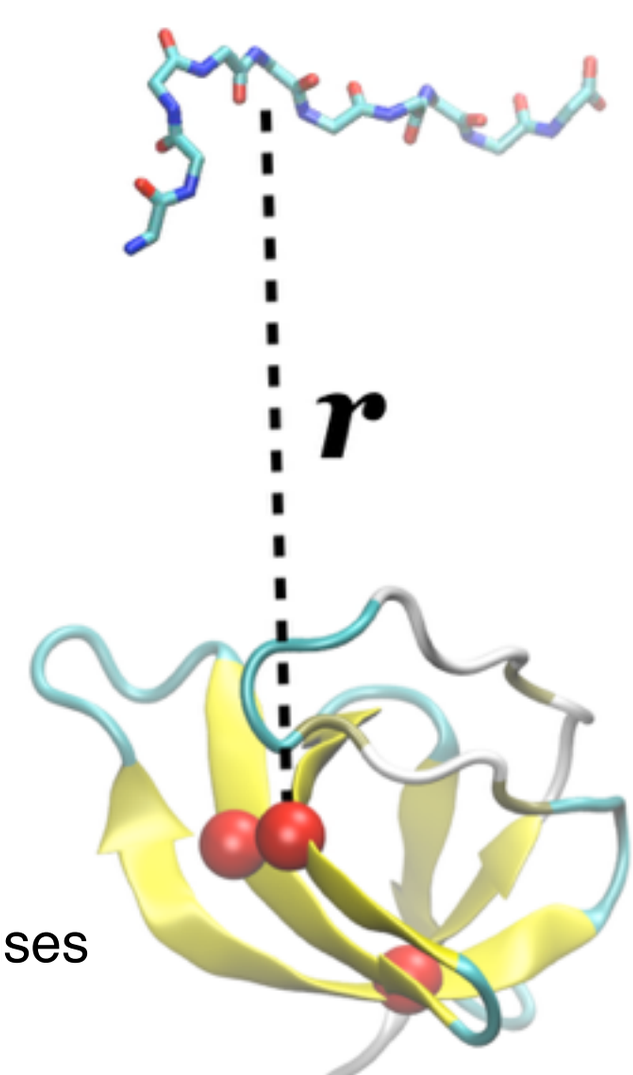

*temperature replica exchange MD works on this principle, but one either uses only a low temperature replica for analysis *OR* unbiases higher *T* replicas

Methods for calculating free energies

(1) histogram-based methods (e.g., umbrella sampling, metadynamics)

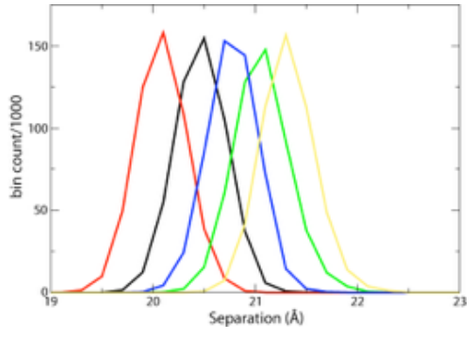

$$\Delta A_i = -\frac{1}{\beta} \ln P_i'(\xi) - w_i(\xi) + F_i$$

biased potentials

(2) non-equilibrium work (e.g., steered MD)

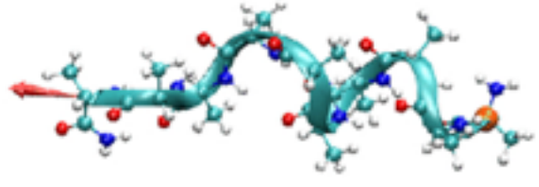

$$\exp(-\beta \Delta A) = \langle \exp(-\beta w) \rangle$$

biased coordinates

(3) alchemical transformations (e.g., free energy perturbation)

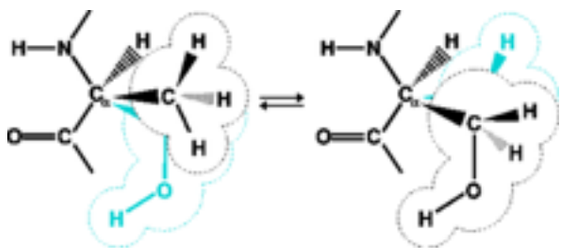

$$\exp(-\beta \Delta A) = \langle \exp(-\beta \Delta U) \rangle_0$$

biased paths

(4) gradient-based methods (e.g., Adaptive Biasing Forces,thermodynamic integration)

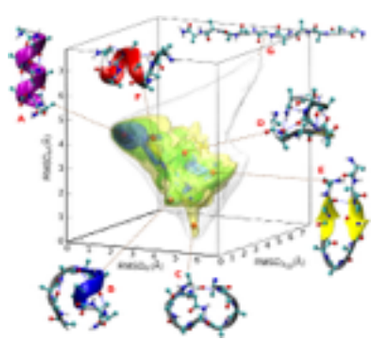

$$\frac{\mathrm{dA}(\xi)}{\mathrm{d}\xi} = \langle \frac{\partial U}{\partial \xi} - \frac{1}{\beta} \frac{\partial \ln |J|}{\partial \xi} \rangle_\xi$$

biased forces

Method 1: Umbrella sampling (US)

Apply restraining potential (bias) on reaction coordinate (RC) for a series of closely spaced windows

$$w_i(\xi) = \frac{1}{2} K (\xi - \xi_i)^2$$

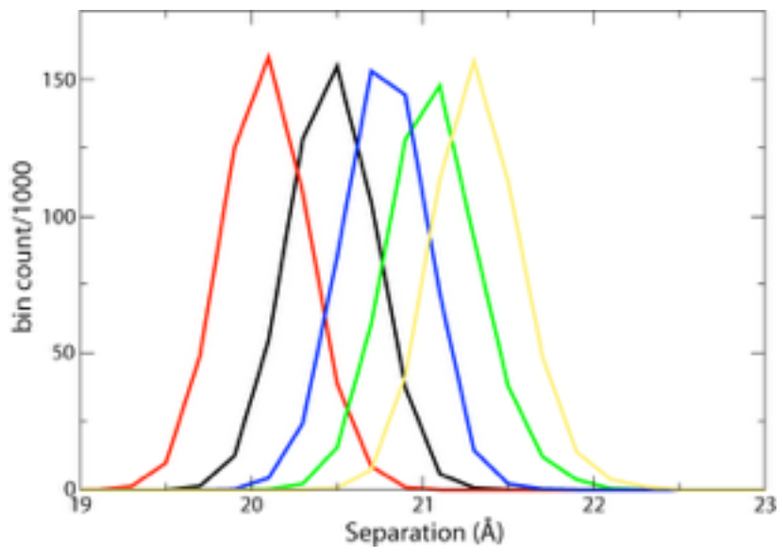

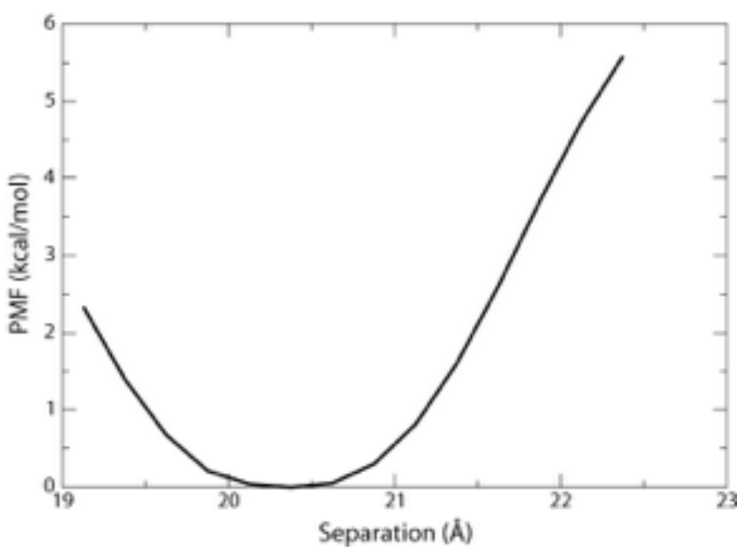

track fluctuations, compute histograms, and unbias/ combine with **WHAM**

$$\Delta A_i = -\frac{1}{\beta} \ln P_i'(\xi) - w_i(\xi) + F_i$$

Grossfield. "WHAM: the weighted histogram analysis method", http://membrane.urmc.rochester.edu/content/wham

replica exchange US: exchanging coordinates periodically between different windows to get around barriers in the RC

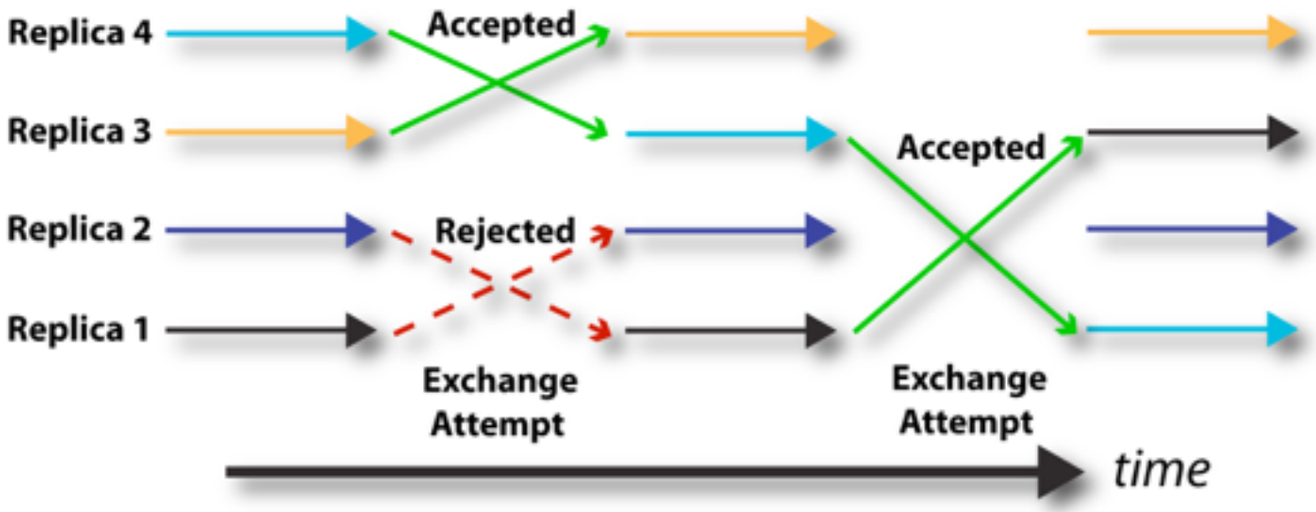

Method 2: Steered molecular dynamics (SMD)

Apply a force to our coordinate of interest to **steer** it in the direction we want it to go

$$\Delta A \leq \langle W \rangle$$

*there is almost always some amount of work wasted**

*except when there is not!

$$e^{-\beta \Delta A} = \langle e^{-\beta W} \rangle$$ **Jarzynski equality**

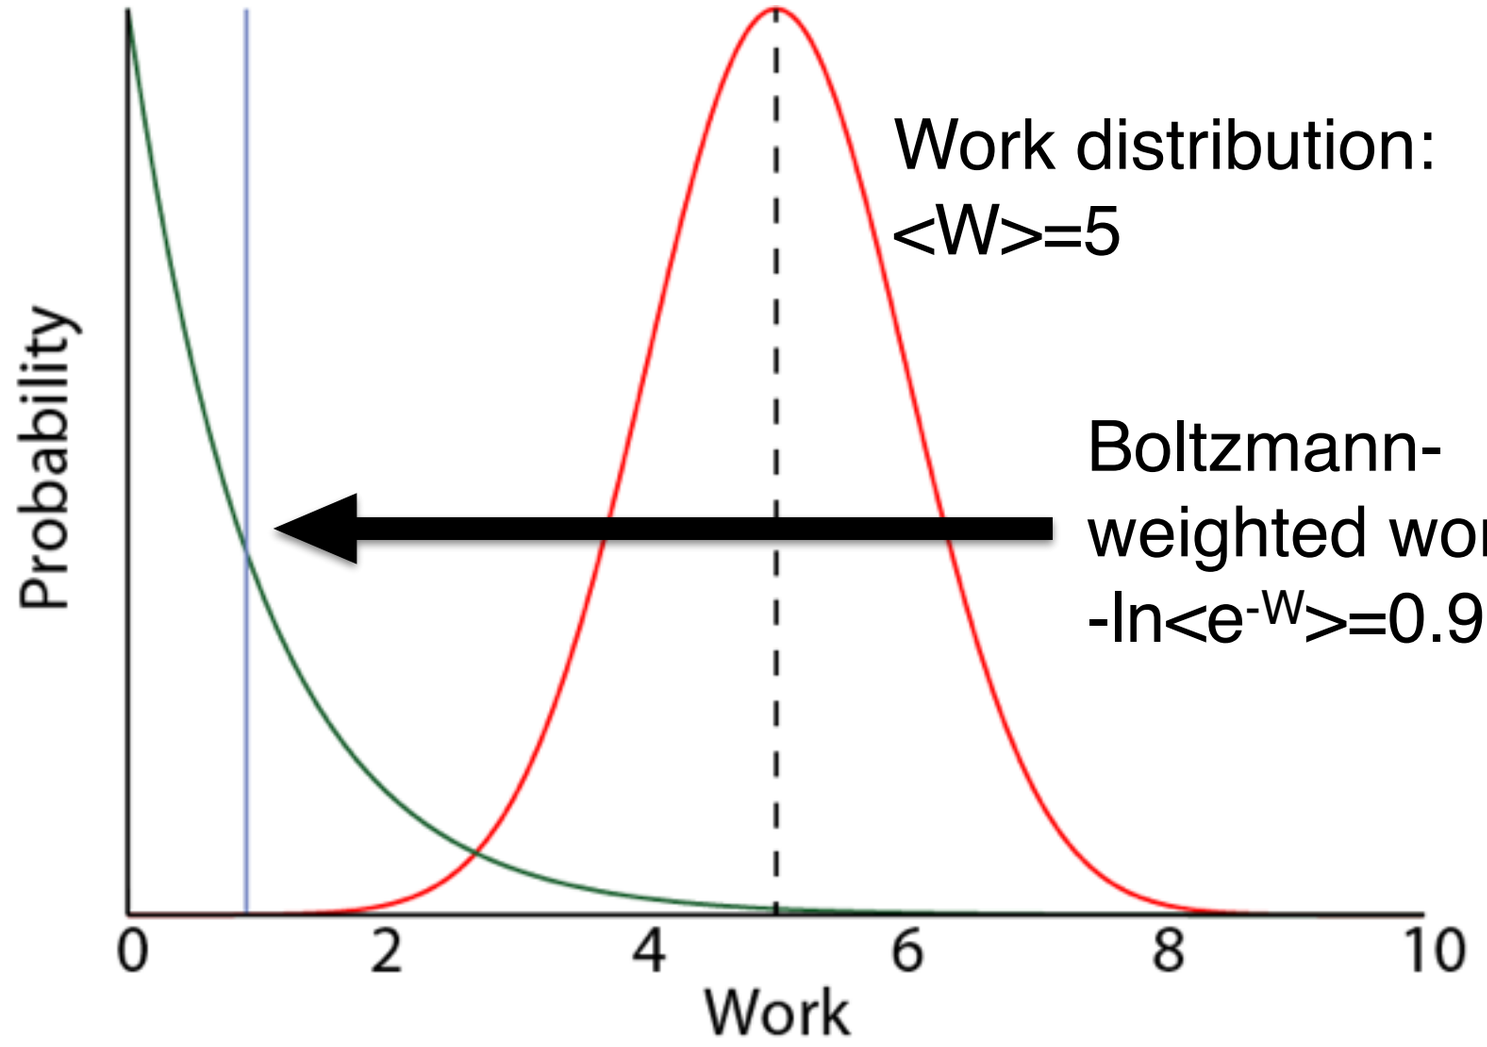

Non-equilibrium methods for free energies are often not practical as they require **significant sampling** to capture **rare events** with spontaneous $W < \Delta A$

Jarzynski. Nonequilibrium equality for free energy differences. (1997) *Phys. Rev. Lett.* 78:2690-2693.

Method 3: Free-energy perturbation (FEP)

Each state (e.g., bound ligand vs. unbound) is represented by its own Hamiltonian

the perturbation

$$H_1(\mathbf{r},\mathbf{p}) = H_0(\mathbf{r},\mathbf{p}) + \Delta H(\mathbf{r},\mathbf{p})$$

Change in free-energy

$$\Delta A = -\frac{1}{\beta}\ln\left(\frac{Z_1}{Z_0}\right)$$

an exercise for the reader!

now expressed as an average in **state 0**

$$\Delta A = -\frac{1}{\beta}\ln\langle\exp[-\beta\Delta H(\mathbf{r},\mathbf{p})]\rangle_0$$

Assume kinetic energy components of *H*s cancel (no change in mass!)

$$\Delta A = -\frac{1}{\beta}\ln\langle\exp[-\beta\Delta U(\mathbf{r})]\rangle_0$$

perturbations need to be small; requires large number of intermediate states

Method 3: Free-energy perturbation (FEP)

$$\Delta A = -\frac{1}{\beta} \ln \langle \exp[-\beta \Delta U(\mathbf{r})] \rangle_0$$

Formally valid for all cases,
but not **practically** so

perturbation approach only converges to the correct answer if there is a strong overlap of phase space ((**x,p**) values accessible with given *H*) in states **0** and **1**, i.e, if the difference between **0** and **1** is small.

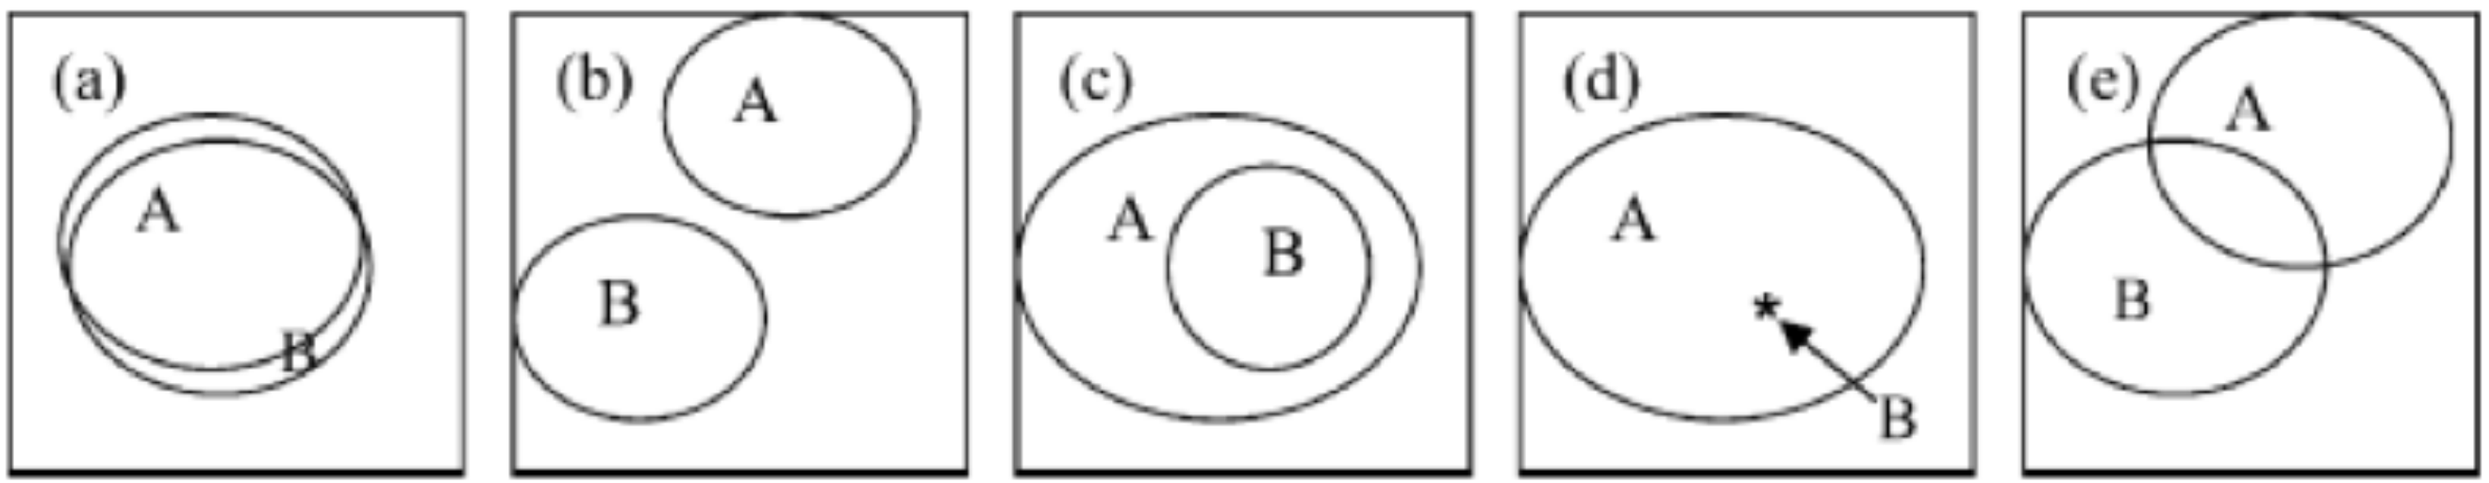

efficient **inefficient** **efficient** **inefficient** **inefficient**

Method 3: Free-energy perturbation (FEP)

in practice, FEP obtains the free energy change by going through a series of **unphysical states**

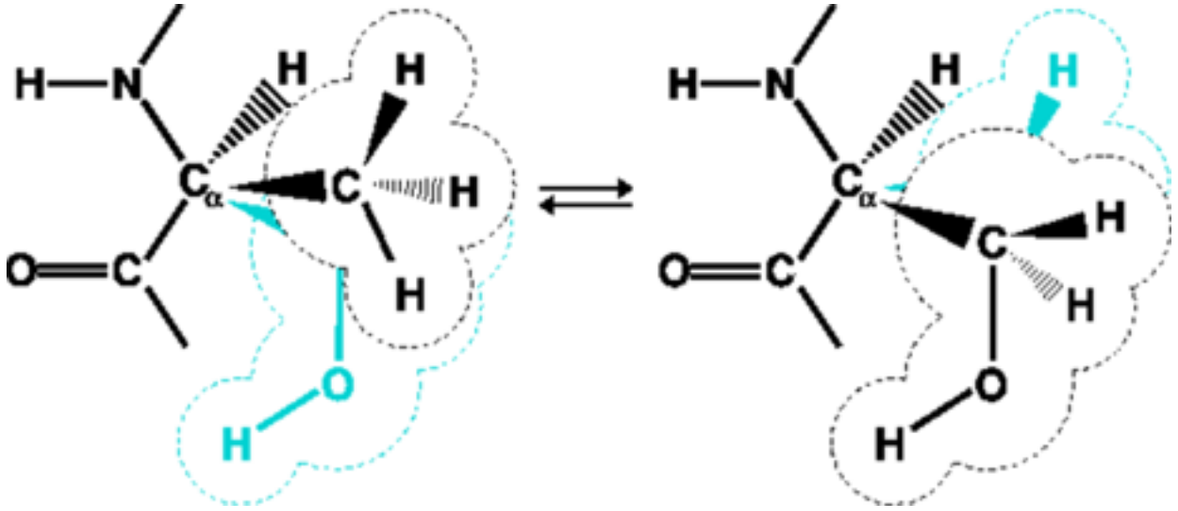

$$H(\mathbf{r},\mathbf{p}) = H_0(\mathbf{r},\mathbf{p}) + \lambda H_b(\mathbf{r},\mathbf{p}) + (1-\lambda)H_a(\mathbf{r},\mathbf{p})$$

-interactions with the environment are scaled by ***λ*** (a coupling parameter) and (**1- *λ***), respectively

Example: both **arginine** and **leucine** sidechains are at the same position on a **polyLeu helix** but don't interact (*dual topology approach*)

leucine/
arginine

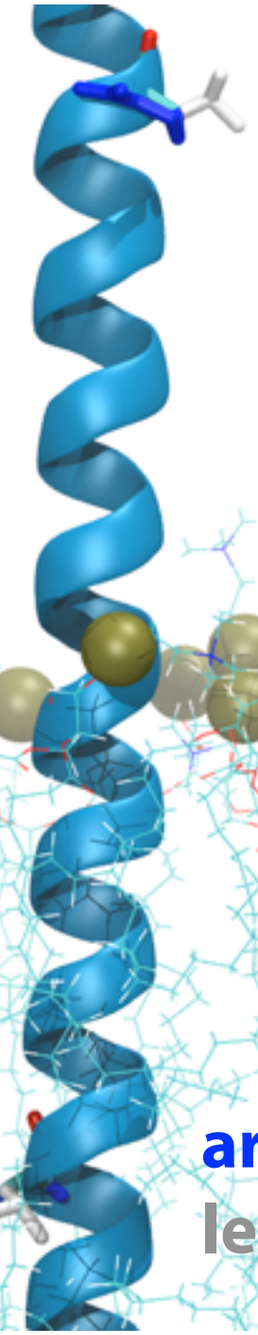

Method 4: Adaptive biasing forces (ABF)

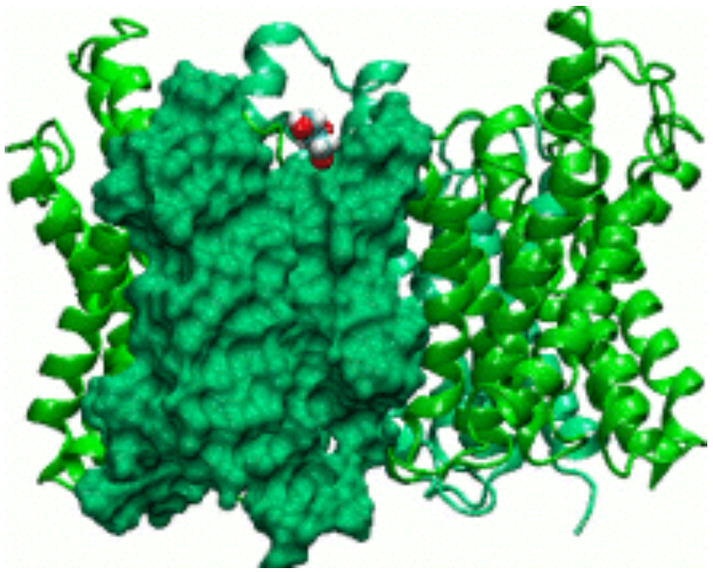

http://www.edam.uhp-nancy.fr/ABF/

$$\Delta A(\xi) = -\frac{1}{\beta} \ln P(\xi)$$

Free energy as function of ξ

$$\nabla_\xi A(\xi) = \langle -F_\xi \rangle_\xi$$

Relation to average force

$$\mathbf{F}^{\mathrm{ABF}} = \nabla_x \tilde{A} = -\langle F_\xi \rangle_\xi \nabla_x \xi$$

Compute average force **adaptively** and apply **biasing force** to cancel it

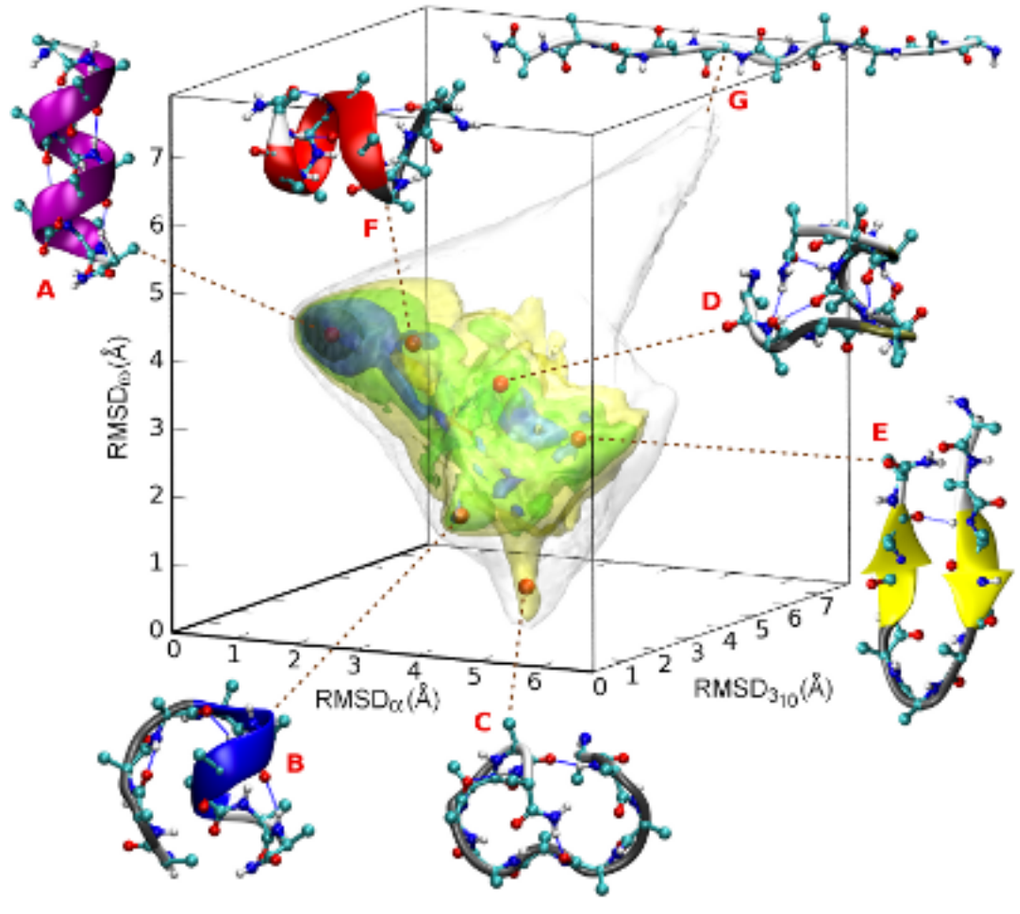

As the estimate of the PMF improves, the biasing forces *should* effectively cancel it, permitting the reaction coordinate to diffuse more easily

Hénin, Fiorin, Chipot, Klein, Exploring Multidimensional Free Energy Landscapes Using Time-Dependent Biases on Collective Variables. (2010) *J. Chem. Theory Comput.*, 6:35-47.

Choosing a *good* Reaction Coordinate

Attempting to reduce a curvilinear path in 3*N*-dimensional space to a **1D path**

this 1D path is known as the **committor probability**; while it exists, it is typically completely non-intuitive — prefer to choose a physically meaningful coordinate that is close to the "true" one

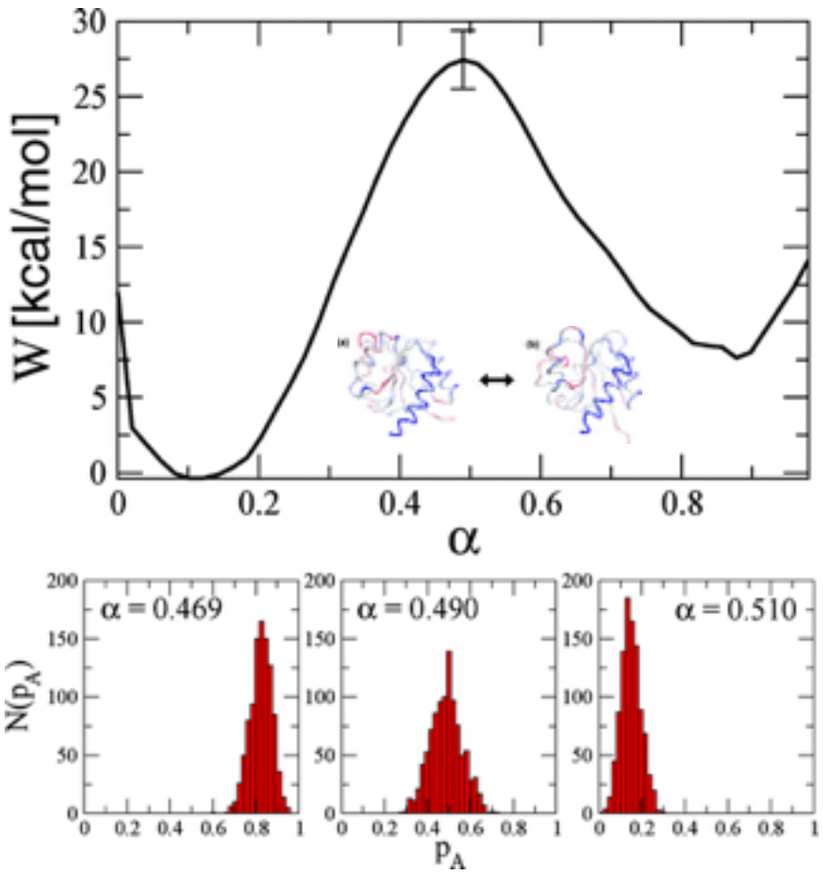

Free energy along a **good** transition coordinate (determined via *string method*)

committor probability is max around **0.5** for the transition state

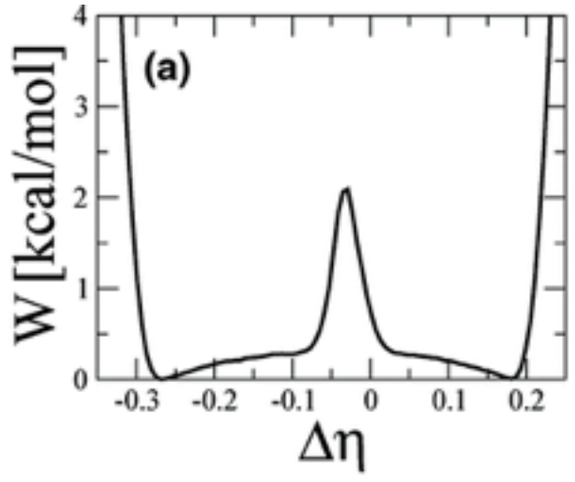

Free energy along a **bad** transition coordinate (fraction of native contacts)

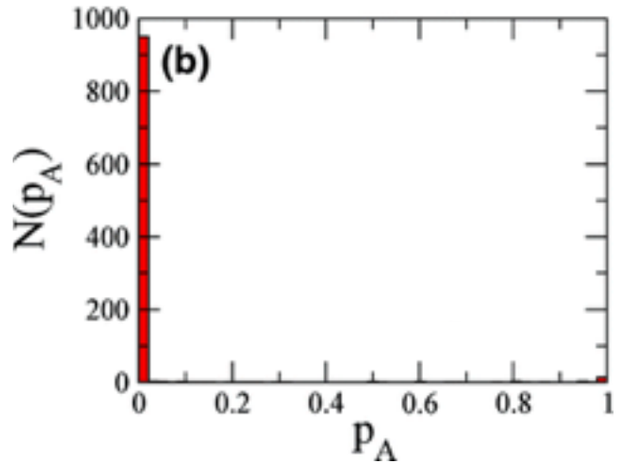

committor probability is all at 0, meaning at the free-energy peak, all configurations fall back to only one min.

Pan, Sezer, Roux. Finding Transition Pathways Using the String Method with Swarms of Trajectories. (2008) *J. Phys. Chem. B*, 112:3432-3440.

Orthogonal degrees of freedom

Because only one (or a few) coordinate(s) is biased in the enhanced sampling simulation, others may yet be slow to evolve

Example:

if ξ is restrained at **0.0** using **US**, a large barrier prevents it from fully sampling the orthogonal ζ

however, if ξ were free to diffuse, the system could take an alternate, lower energy path to reach **state 2** from **state 1** faster

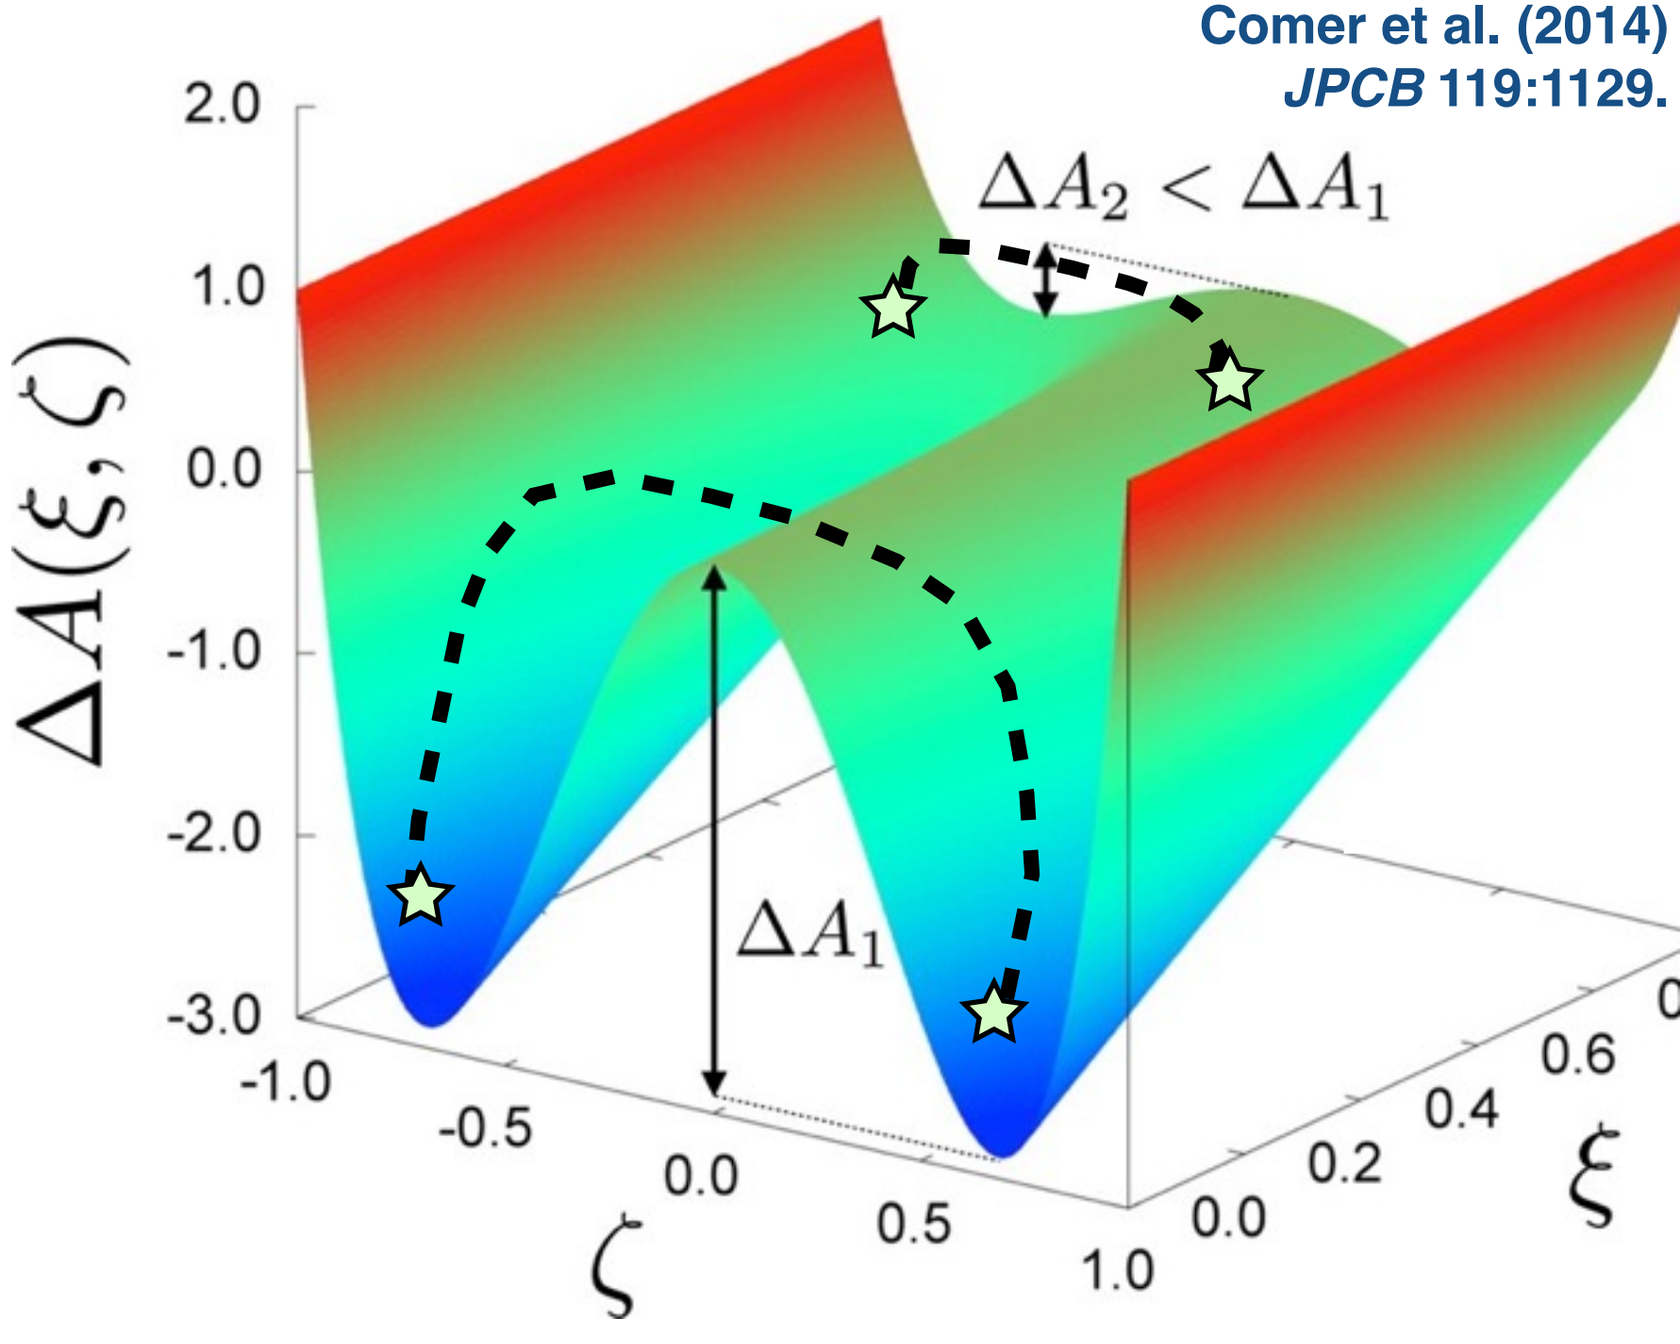

ALL methods suffer from a version of this sampling difficulty!

Two forms of error and its estimation

1) statistical error (the known unknowns)

every free-energy method has ways of estimating this, although it is typically small and can be reduced through increased sampling

Ex: variance in free energy as a function of variance in each window for Umbrella Sampling simulation

$$\mathrm{var}[G(\xi)] \approx (K\Delta\xi)^2 \cdot \sum_{i=1}^{(\xi-\xi_0)/\Delta\xi} \mathrm{var}(\bar{\xi}_i)$$

Zhu, Hummer. Convergence and error estimation in free energy calculations using the weighted histogram analysis method. (2012) *J. Comput. Chem.* 33:453–465.

$$\sigma^2_{\Delta A} = \frac{1}{N_0\beta^2}\left[\frac{\langle f^2(x)\rangle_0}{\langle f(x)\rangle_0^2} - 1\right] + \frac{1}{N_1\beta^2}\left[\frac{\langle f^2(-x)\rangle_1}{\langle f(-x)\rangle_1^2} - 1\right]$$

$$x = \beta(\Delta U - C)$$

$$f(x) = \frac{1}{1+\exp(x)}$$

*Ex: variance in stratified FEP simulations using the Bennett Acceptance Ratio (**BAR**) estimator*

Pohorille, Jarzynski, Chipot. Good practices in free-energy calculations. (2010) *J. Phys. Chem. B.* 114:10235-53.

a common way of estimating statistical error is to use **block averaging**, where the simulation is divided into uncorrelated blocks, the free energy determined for each, and then the **standard error of means** is calculated

Two forms of error and its estimation

2) systematic error (the unknown unknowns)

much more difficult to estimate; sources include force field inaccuracies, poor overlap of neighboring windows, and quasi non-ergodic scenarios (e.g., sampling only one of two metastable states)

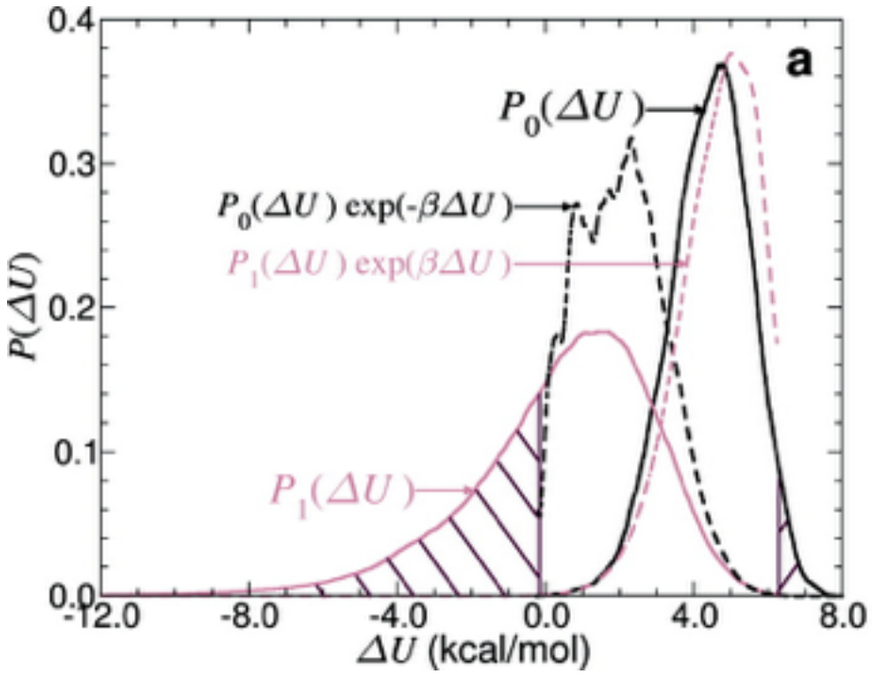

poor overlap

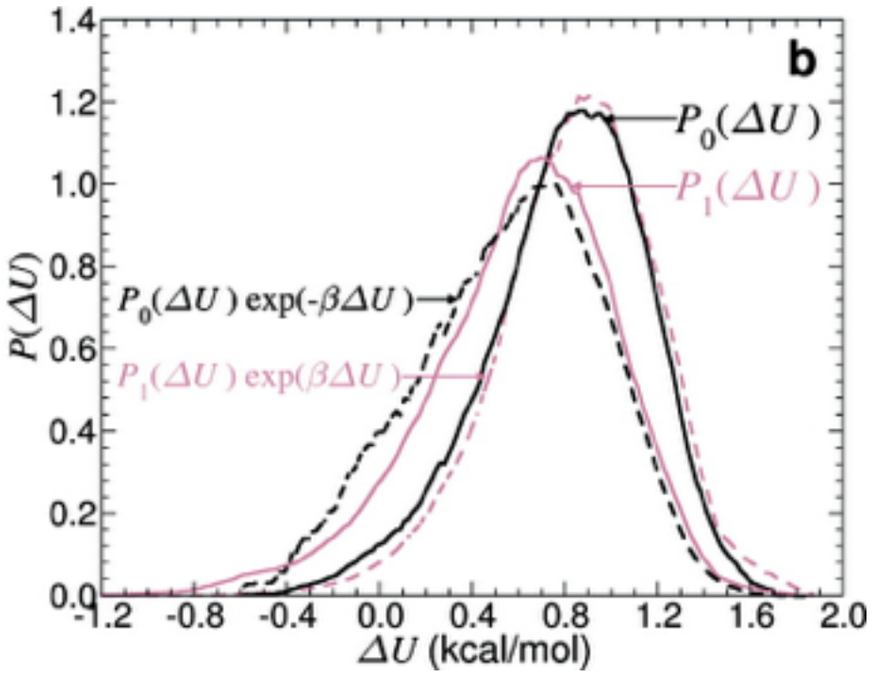

good overlap

*overlap between forward and backward **FEP** runs, **NOT** neighboring windows in one run*

typical ways of estimating this error are to look at overlap between windows (**umbrella sampling, FEP**) or continuity of forces between windows (**ABF**)

Comer…Chipot. The Adaptive Biasing Force method: Everything you always wanted to know but were afraid to ask. (2010) *J. Phys. Chem. B.* 119:1129–1151.

Pohorille, Jarzynski, Chipot. Good practices in free-energy calculations. (2010) *J. Phys. Chem. B.* 114:10235-53.

Example 1: Protein folding

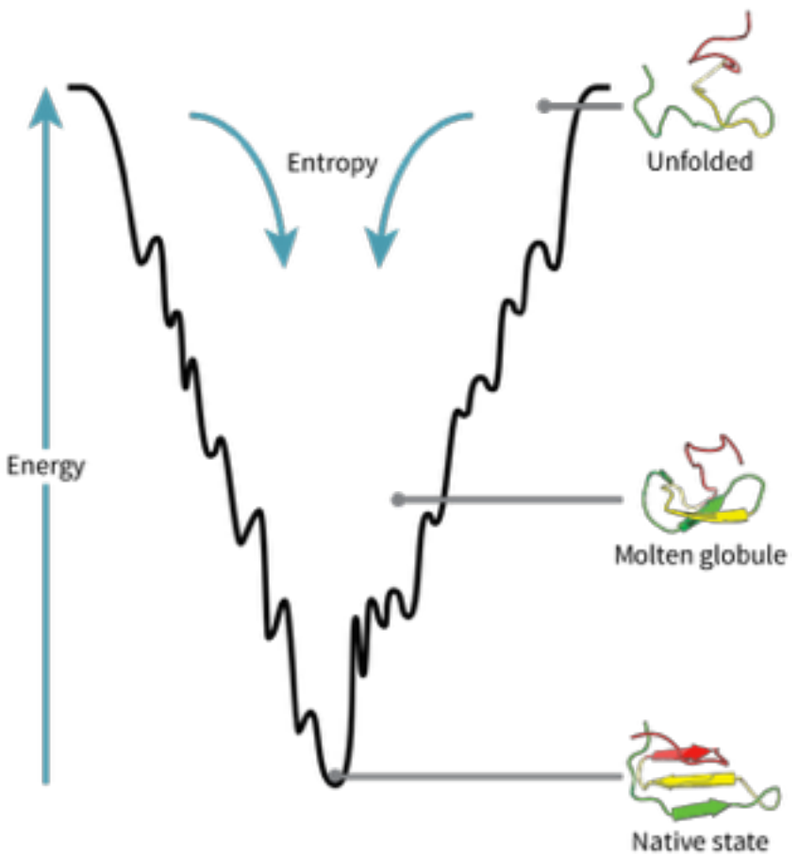

folding funnel

proceeds through a series of intermediate states

true free-energy landscape is typically much more complex

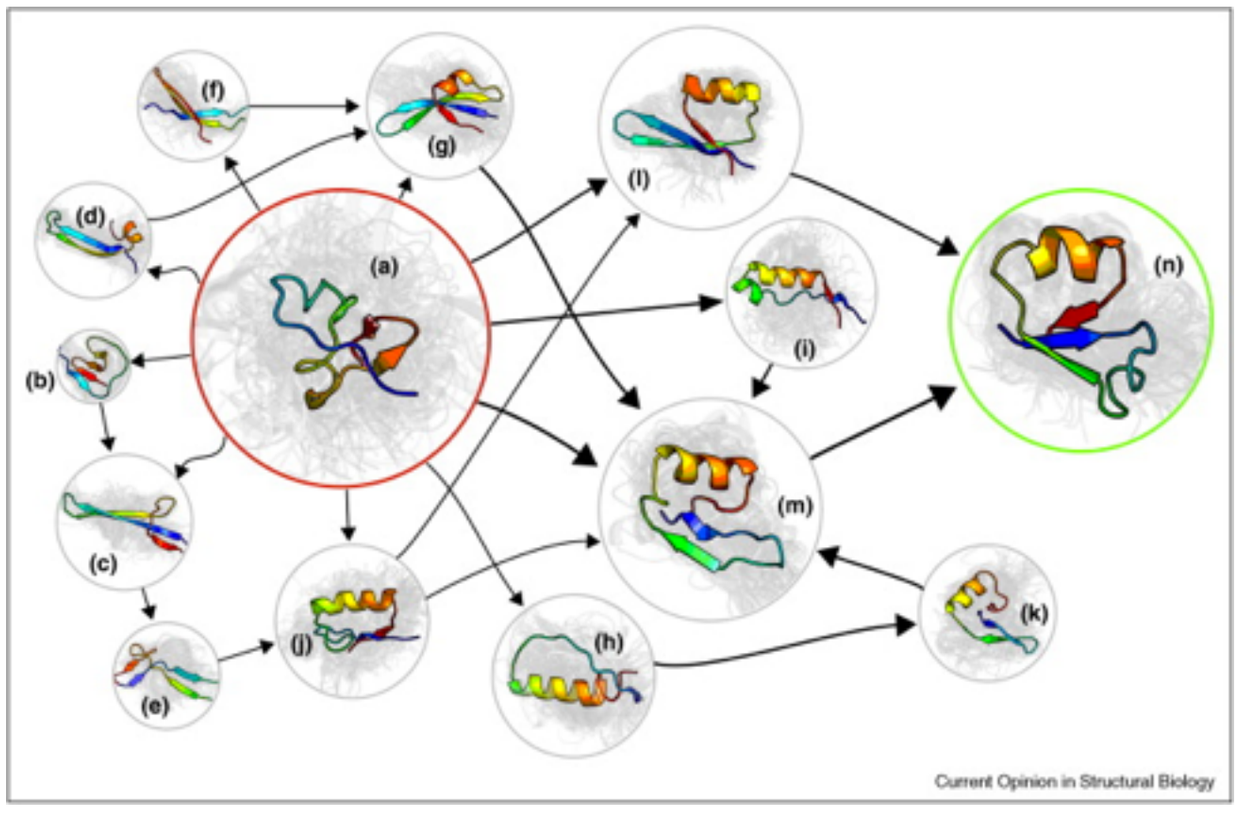

F Ulrich Hartl & Manajit Hayer-Hartl. Converging concepts of protein folding in vitro and in vivo Nat. Struct. Mol. Bio. 16, 574 - 581 (2009)

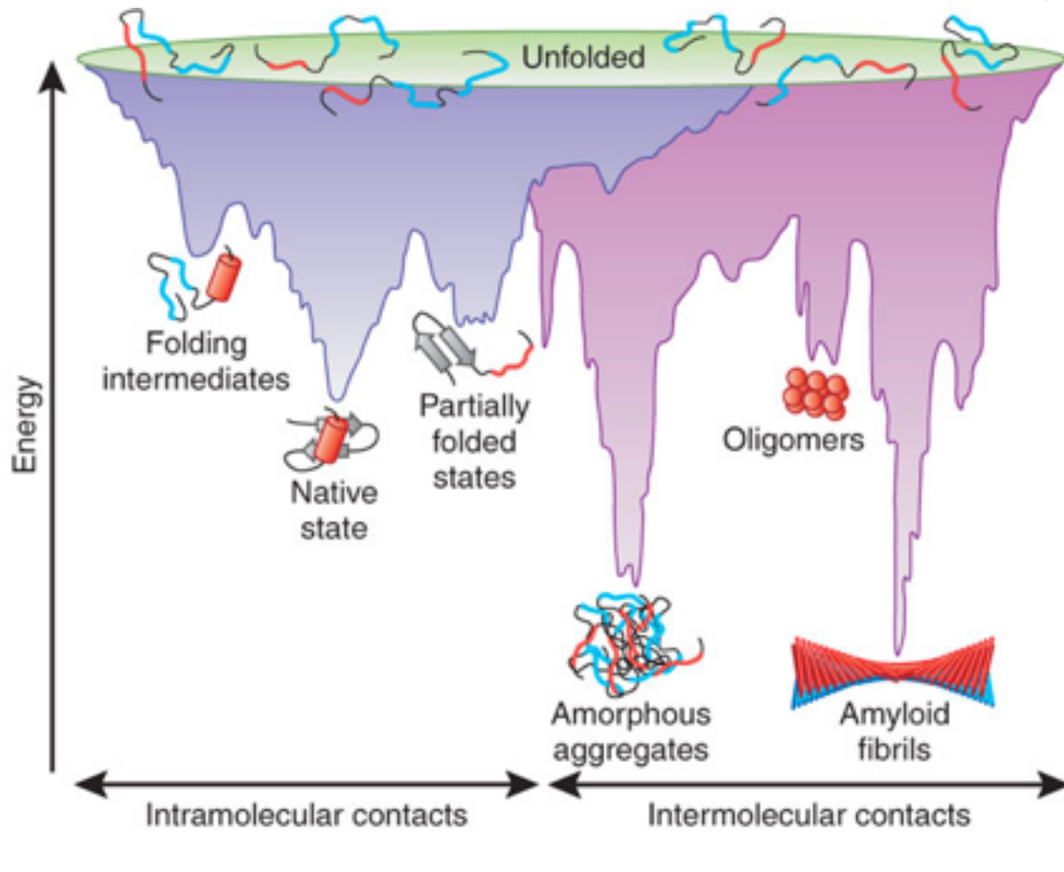

proteins may sample a number of intermediates *without native-like structure* on the folding pathway - *what is the reaction coordinate?*

Bowman, G. R.; Voelz, V. A.; Pande, V. S.Taming the complexity of protein folding Curr. Opin. Struct. Biol. 2011, 21, 4– 11

deca-alanine folding

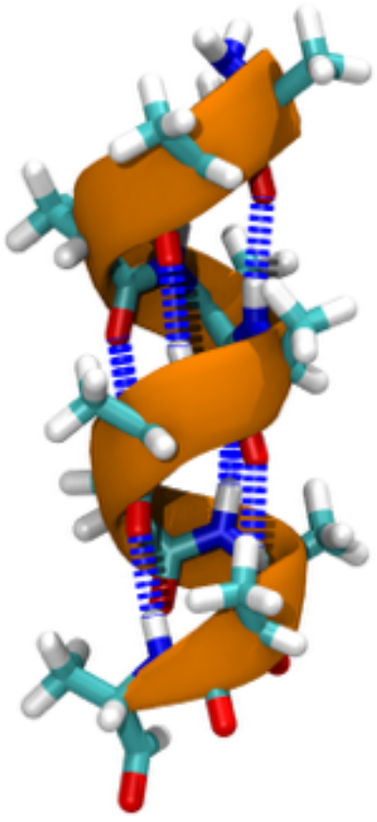

14 Å

32 Å

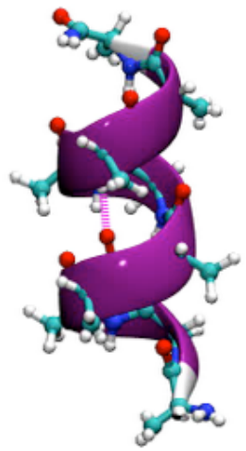

10-Ala helix (in vacuum)
end-to-end distance (*ξ*)
a common **RC**

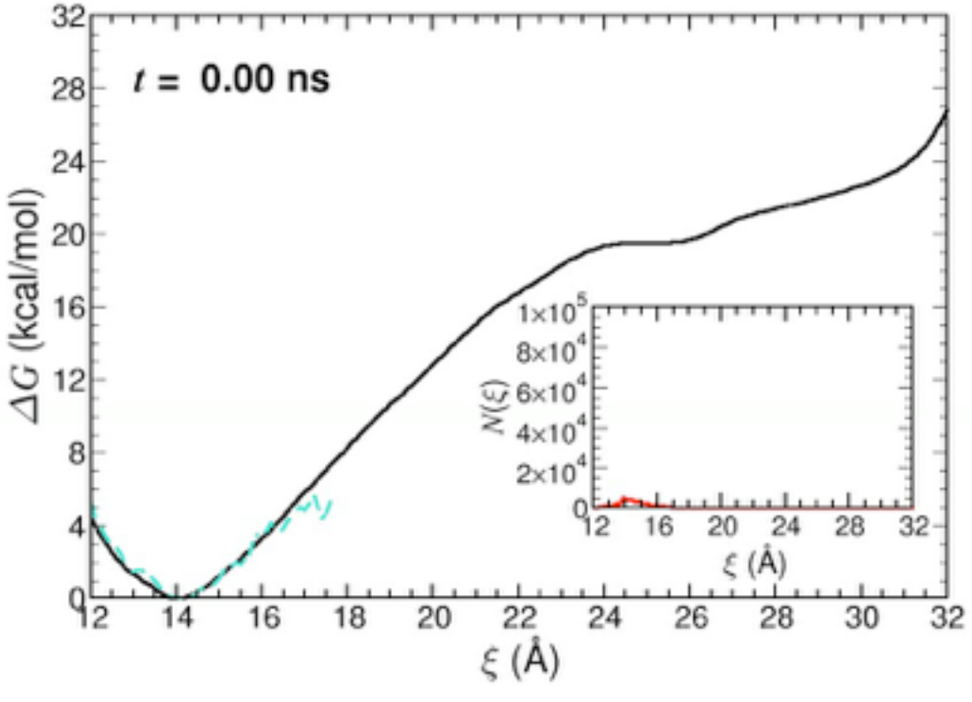

Chipot, Hénin. Exploring the free energy landscape of a short peptide using an average force. (2005) *J. Chem. Phys.* 123:244906.

folding Ala_{10} in vacuum

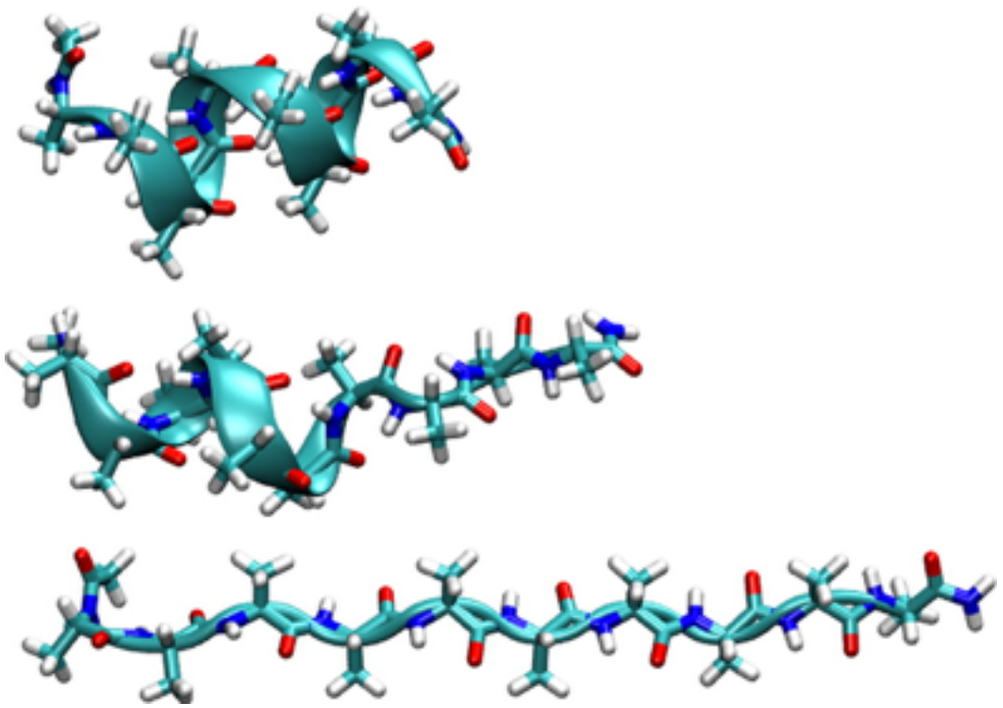

Ala_{10} folds/unfolds in an accordion-like fashion in vacuum

end-to-end distance works well as an RC for both methods

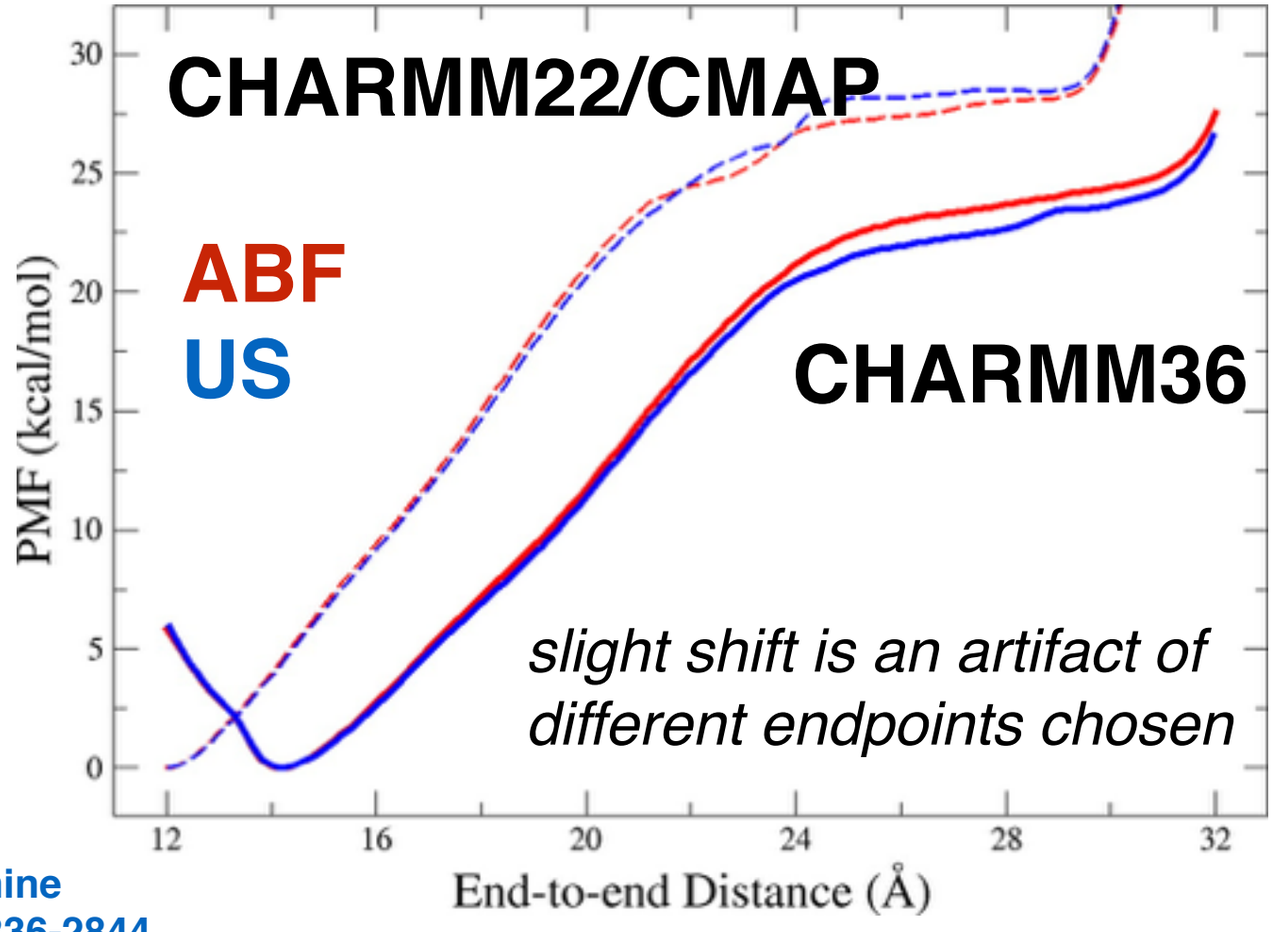

Hazel, Chipot, Gumbart. Thermodynamics of deca-alanine folding in water. (2014) *J. Chem. Theory Comput.* 10: 2836-2844.

folding Ala_{10} in water

end-to-end distance no longer works - **PMFs never converge**!

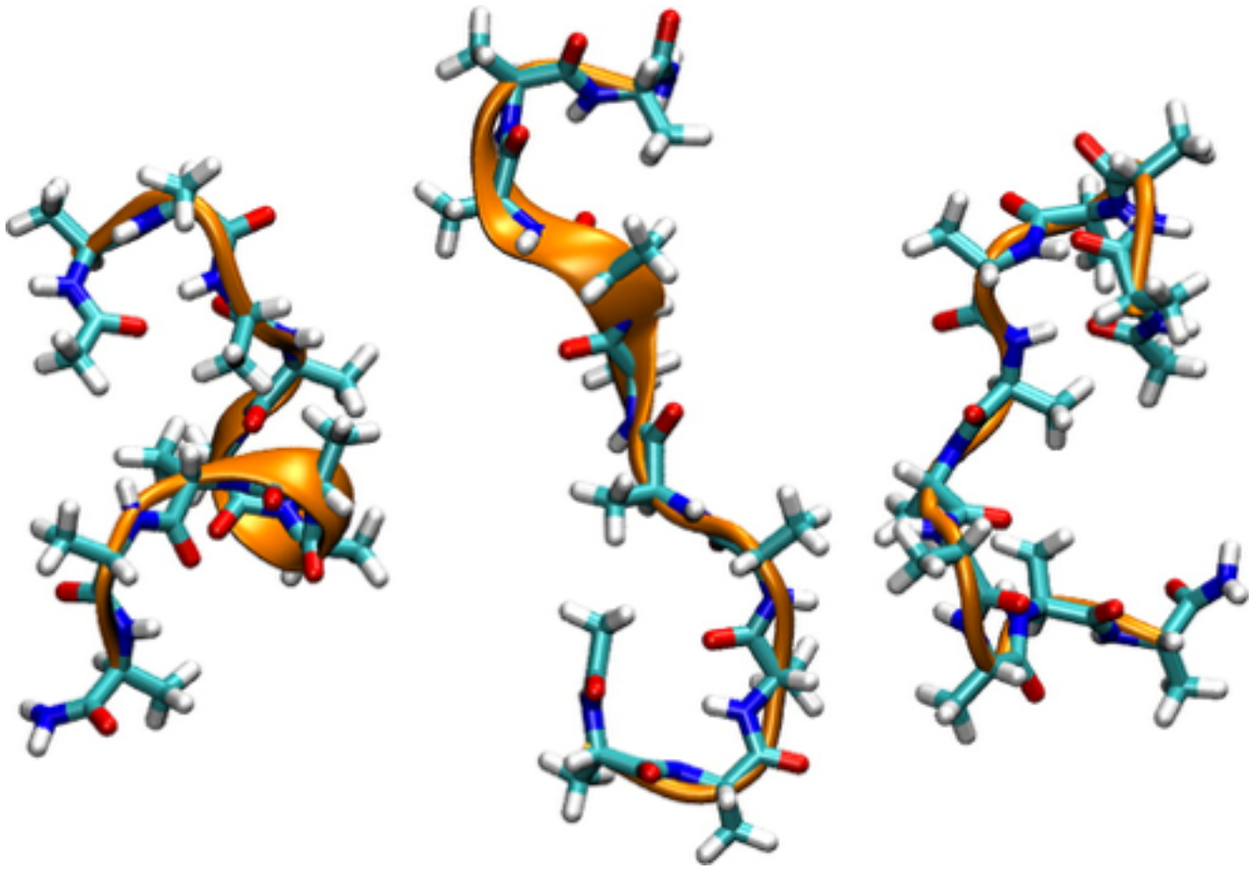

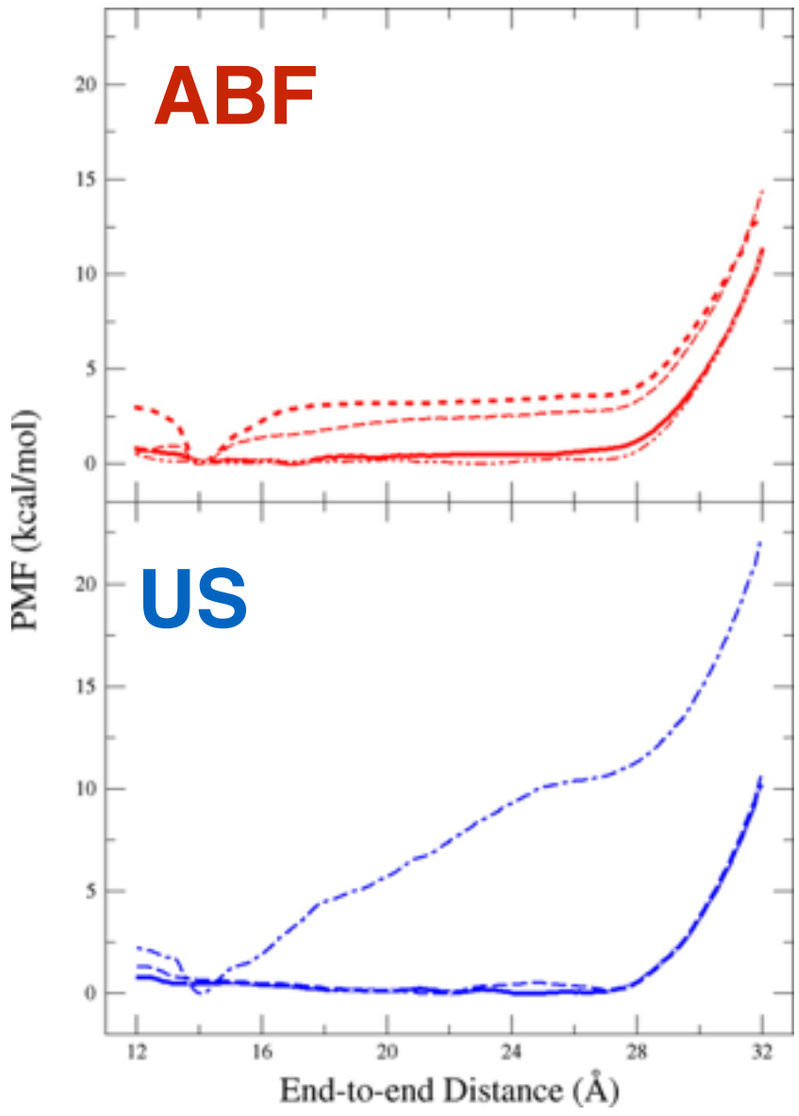

numerous states not observed in vacuum appear

restraints restricting conformational freedom tried to no avail

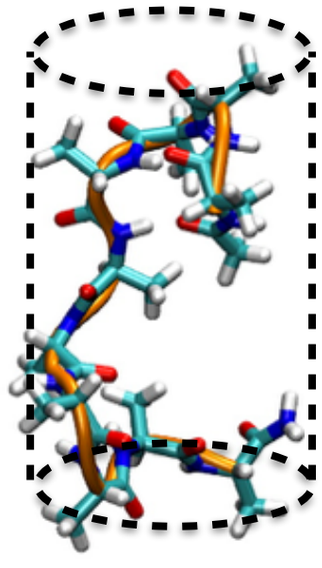

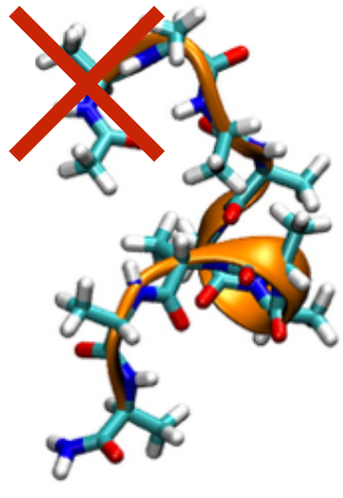

Going to a 2D description

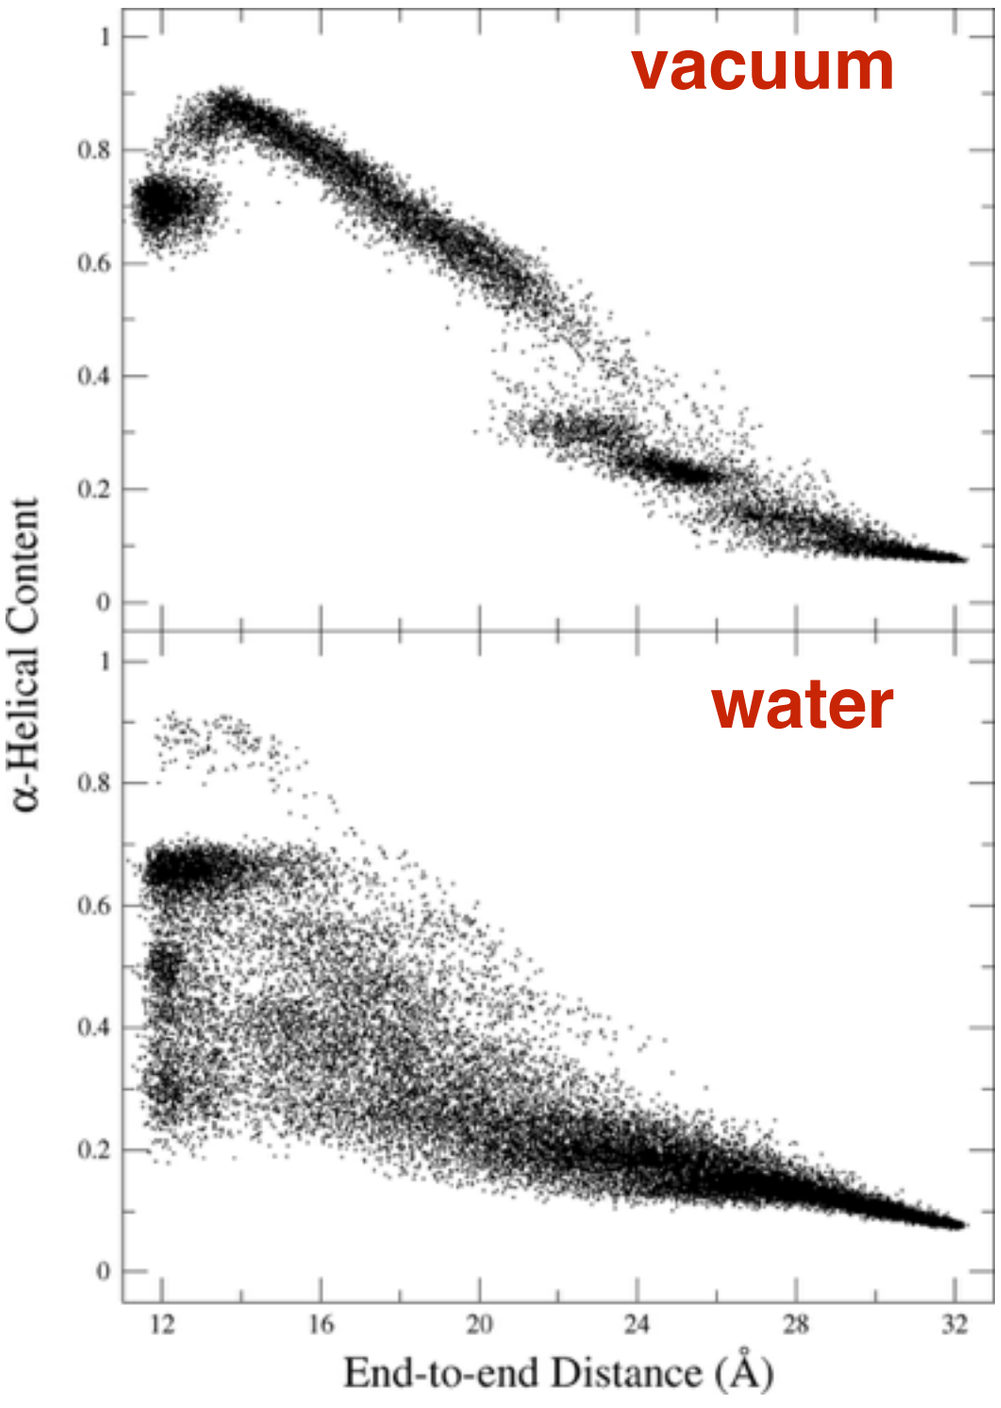

in **vacuum**, alpha helical content and end-to-end distance are practically 1-1

in **water**, a number of compact, low-lying states appear that "contaminate" the 1D PMF (i.e., *are poorly sampled*)

Solution: add a 2nd RC for alpha helicity (***α***), in addition to ***ξ***

$$\alpha\left(C_\alpha^{(N_0)}, O^{(N_0)}, C_\alpha^{(N_0+1)}, O^{(N_0+1)}, \ldots N^{(N_0+5)}, C_\alpha^{(N_0+5)}, O^{(N_0+5)}, \ldots N^{(N_0+N)}, C_\alpha^{(N_0+N)}\right) = \frac{1}{2(N-2)} \sum_{n=N_0}^{N_0+N-2} \text{angf}\left(C_\alpha^{(n)}, C_\alpha^{(n+1)}, C_\alpha^{(n+2)}\right) + \frac{1}{2(N-4)} \sum_{n=N_0}^{N_0+N-4} \text{hbf}\left(O^{(n)}, N^{(n+4)}\right)$$

$$\text{angf}\left(C_\alpha^{(n)}, C_\alpha^{(n+1)}, C_\alpha^{(n+2)}\right) = \frac{1 - \left(\left|\theta\left(C_\alpha^{(n)}, C_\alpha^{(n+1)}, C_\alpha^{(n+2)}\right) - \theta_0\right| / \Delta\theta_{tol}\right)^2}{1 - \left(\left|\theta\left(C_\alpha^{(n)}, C_\alpha^{(n+1)}, C_\alpha^{(n+2)}\right) - \theta_0\right| / \Delta\theta_{tol}\right)^4}$$

$$\text{hbf}\left(O^{(n)}, N^{(n+4)}\right) = \frac{1 - \left(\left|\mathbf{x}_{O^{(n)}} - \mathbf{x}_{N^{(n+4)}}\right| / d_0\right)^n}{1 - \left(\left|\mathbf{x}_{O^{(n)}} - \mathbf{x}_{N^{(n+4)}}\right| / d_0\right)^m}$$

Going to a 2D description

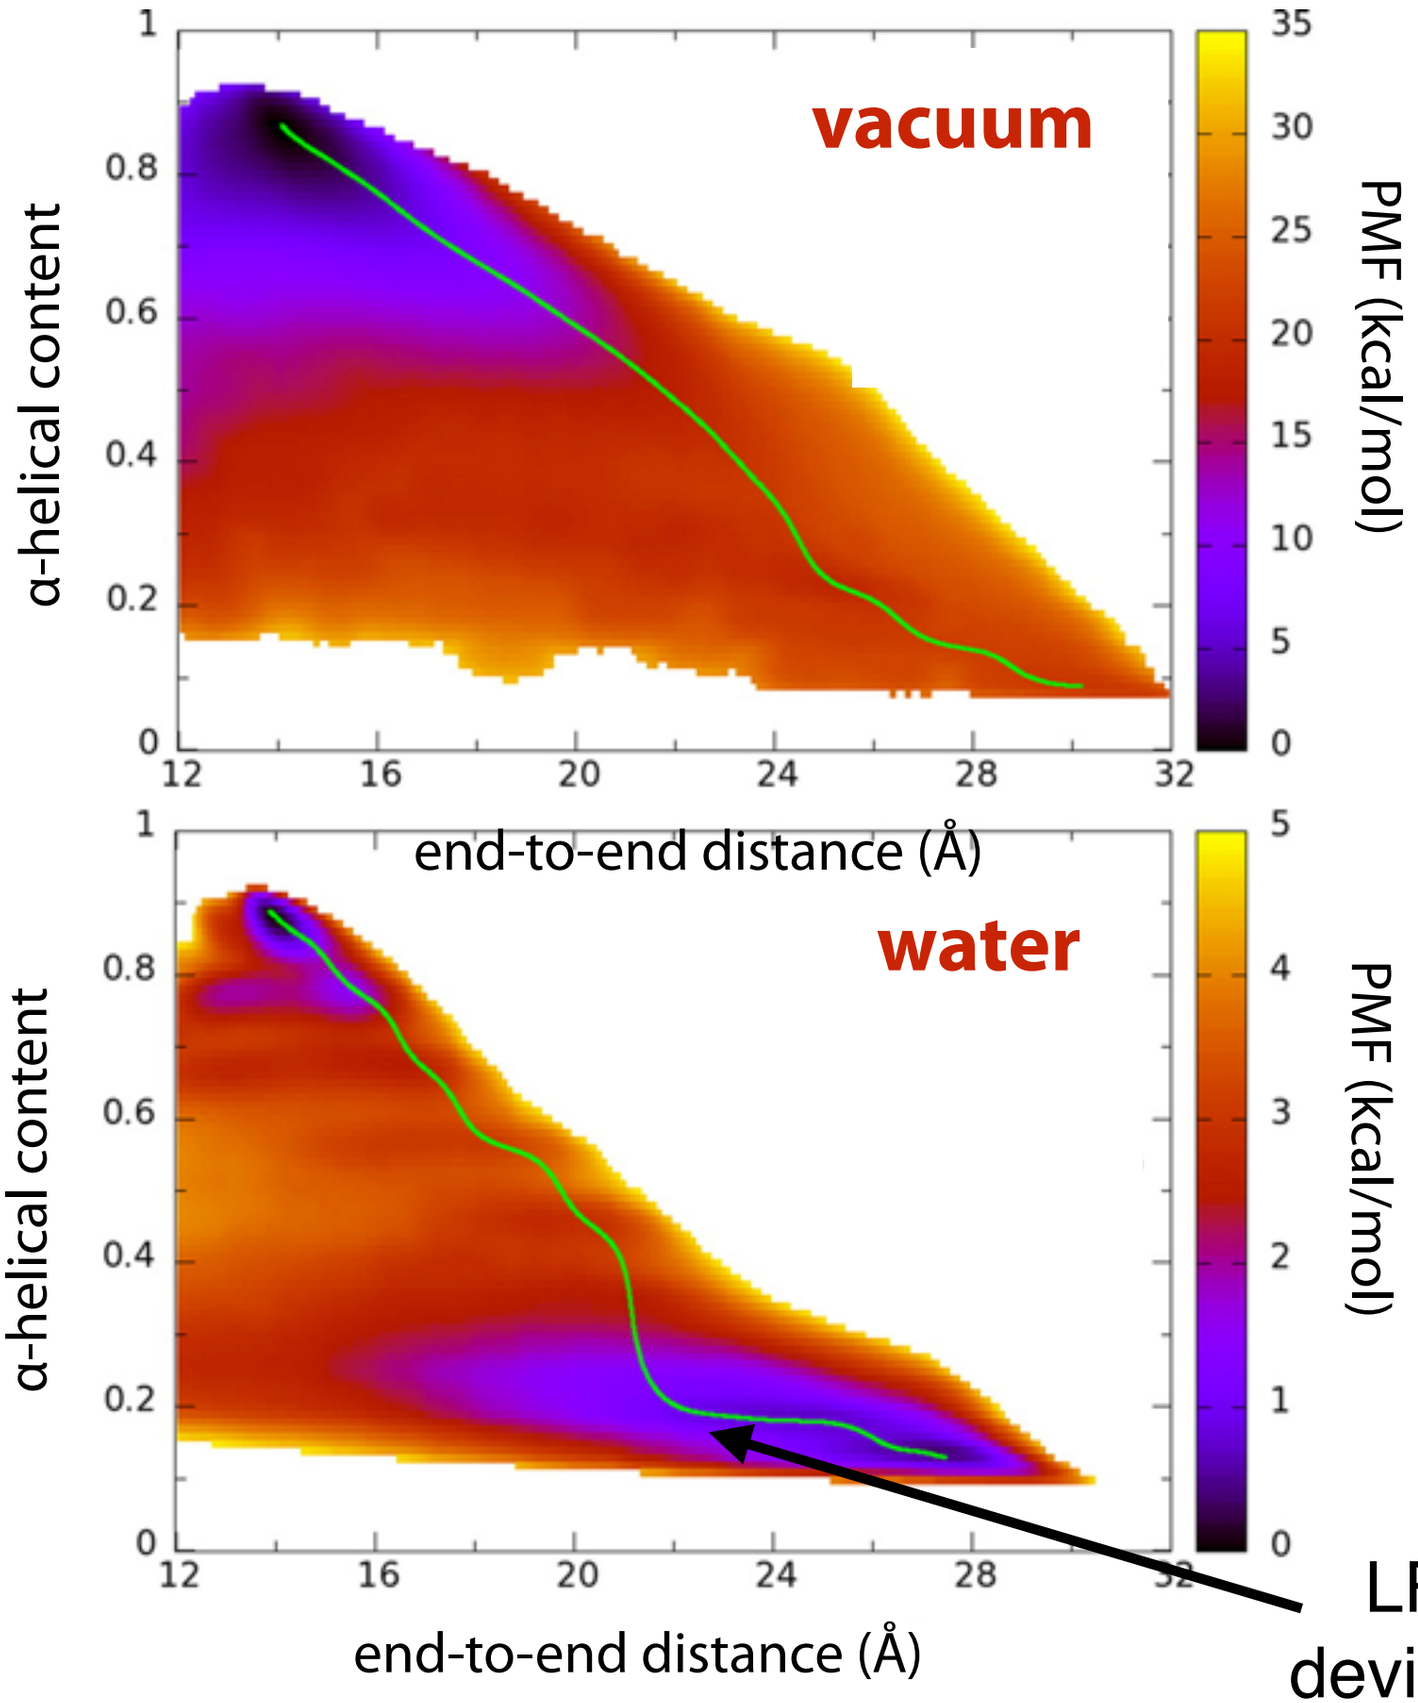

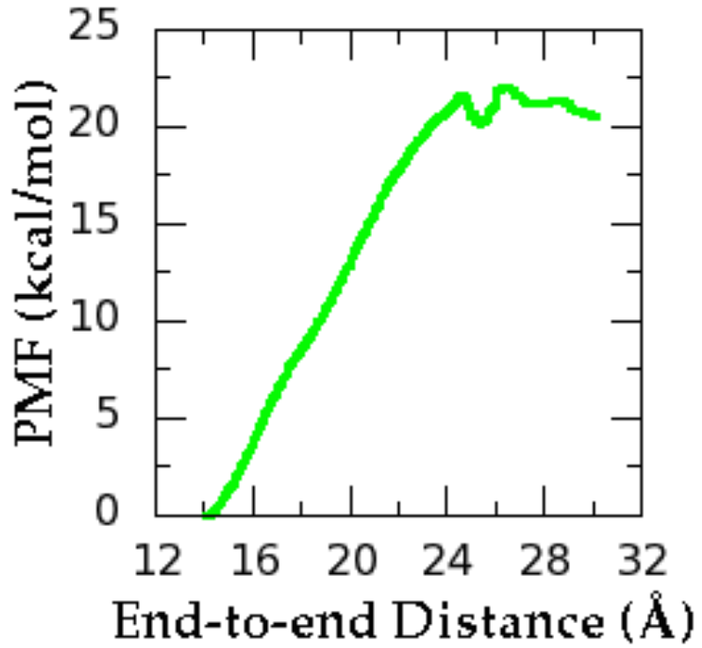

Least free-energy path (**LFEP***) agrees well with 1D vacuum PMF

*Ensing, Laio, Parrinello, Klein. (2005) *J. Phys. Chem. B.* 109:6676– 6687.

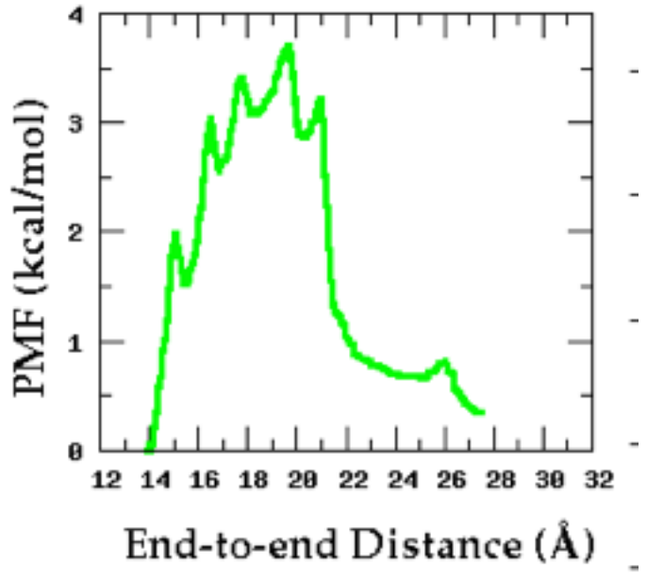

LFEP from 2D PMF in water shows deviation from “accordion-like” unfolding

Is the LFEP the real pathway?

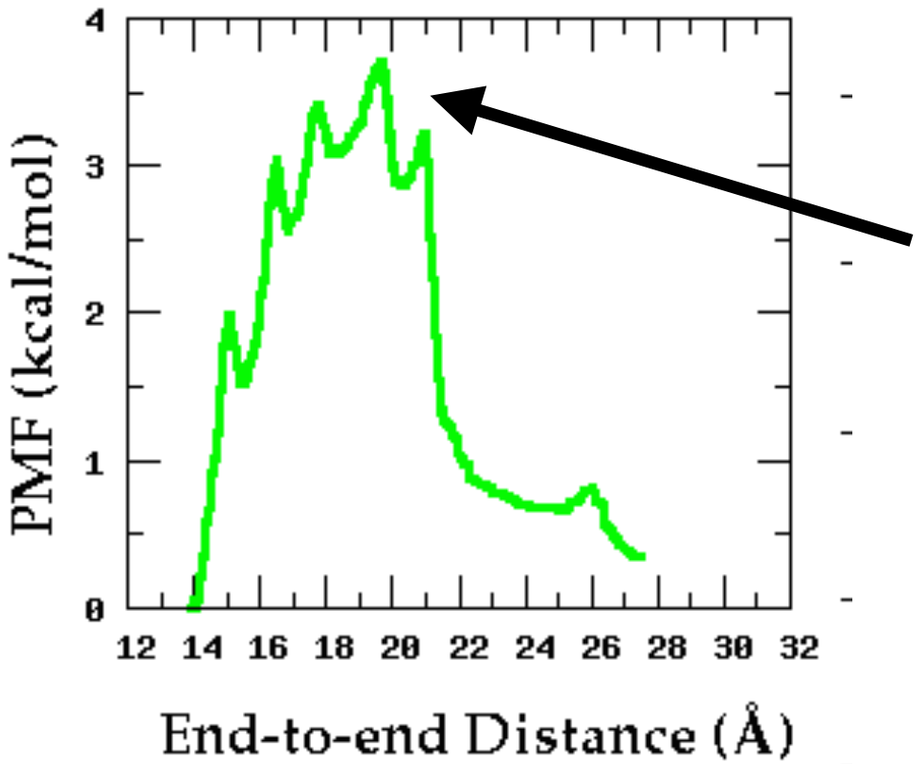

If the choice of RCs is reasonable, the peptide should have a **50/50 probability** of going either way at the free-energy maximum

estimate of the **committor probability distribution** is roughly peaked around 0.5 → a true transition state

(~100 conformations * 50 10-ps simulations each)

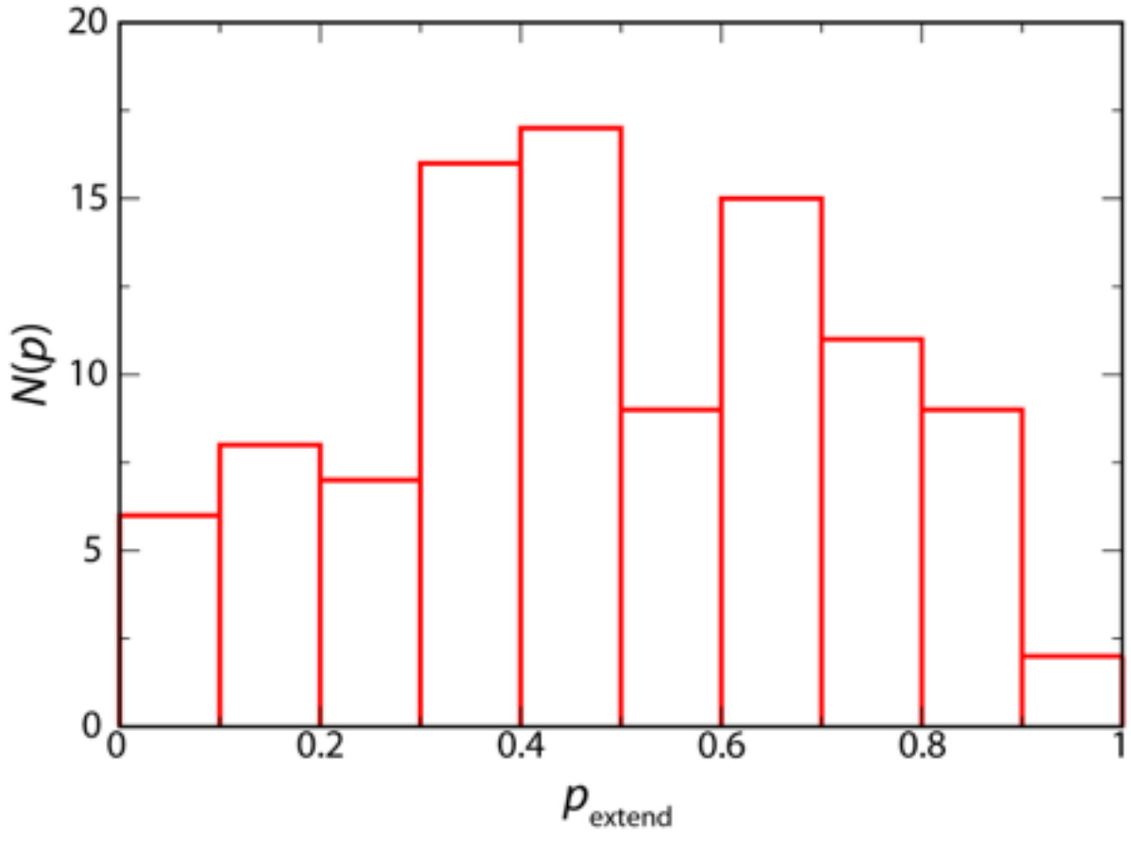

Revisiting the 1D PMF

Integrate out ***α*** RC to get 1D PMF as a function of distance ***ξ***, as originally intended

compact states only ***very slightly favored*** over extended; barrier between them is ~2kT

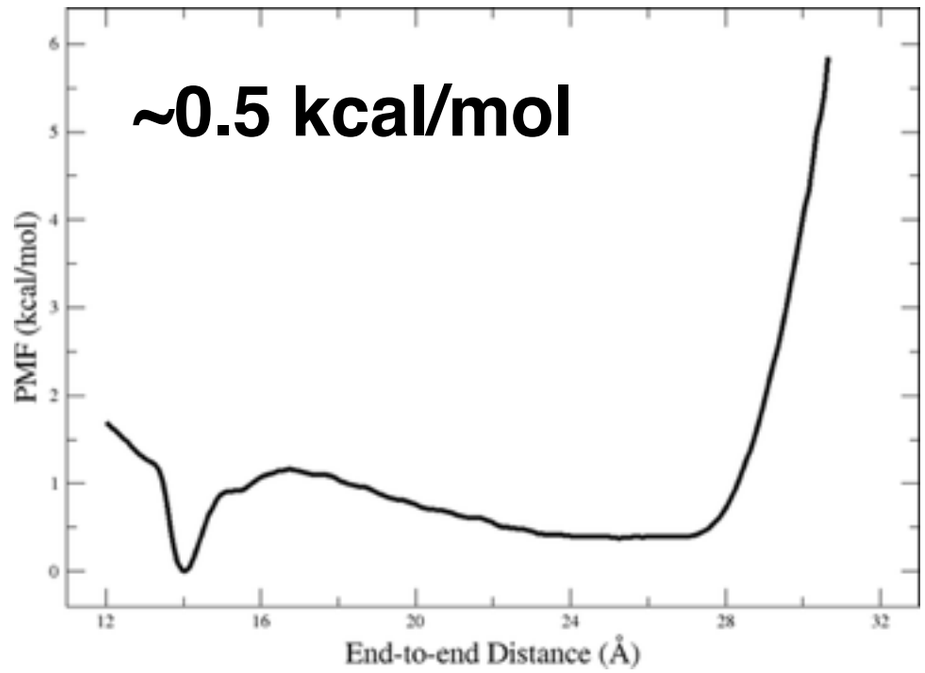

Comparison to other published PMFs

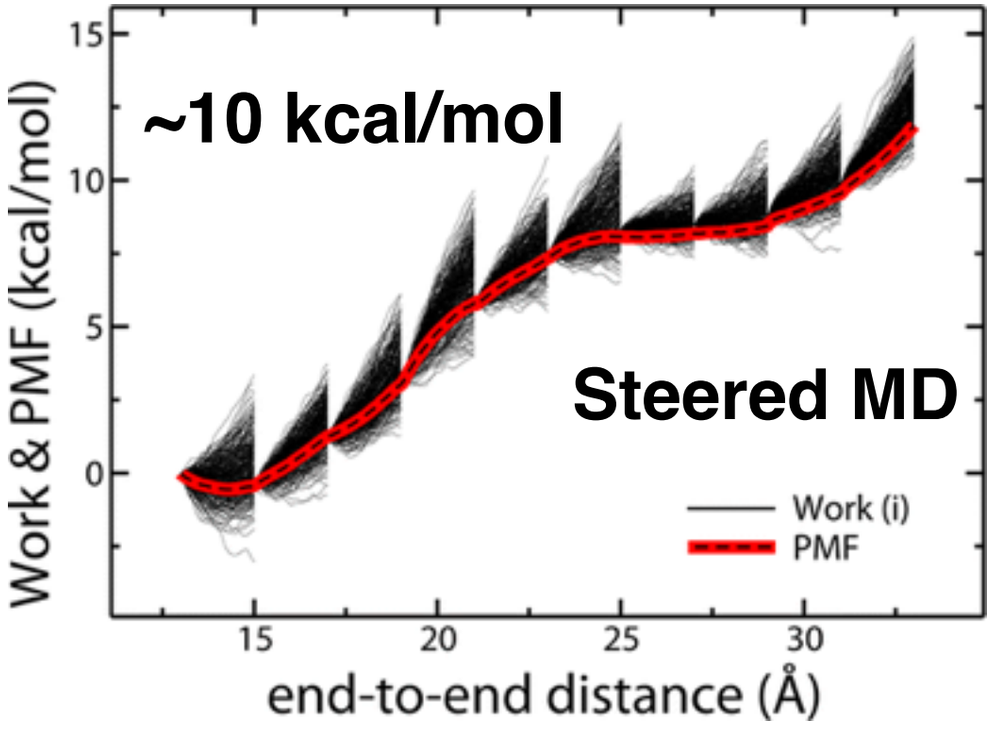

Ozer, Quirk, Hernandez.Thermodynamics of decaalanine stretching in water obtained by adaptive steered molecular dynamics simulations. (2012) *J. Chem. Theory Comput.* 8:4837– 4844.

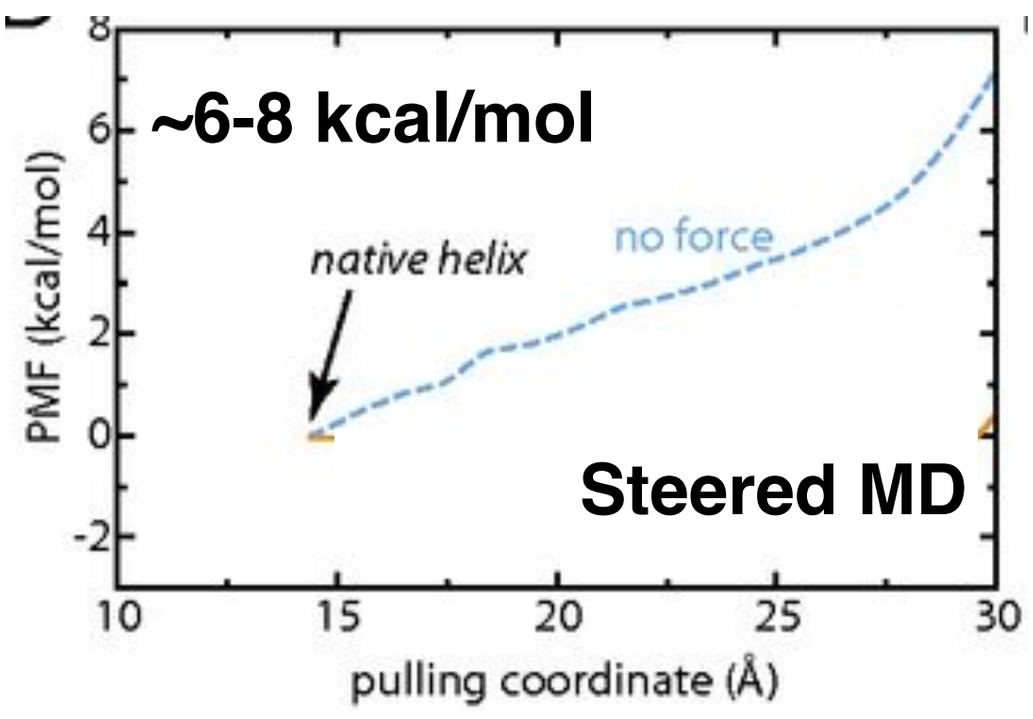

Stirnemann, Kang, Zhou, Berne. How force unfolding differs from chemical denaturation. (2014) *PNAS* 111:3413-3418.

Do we trust our results?

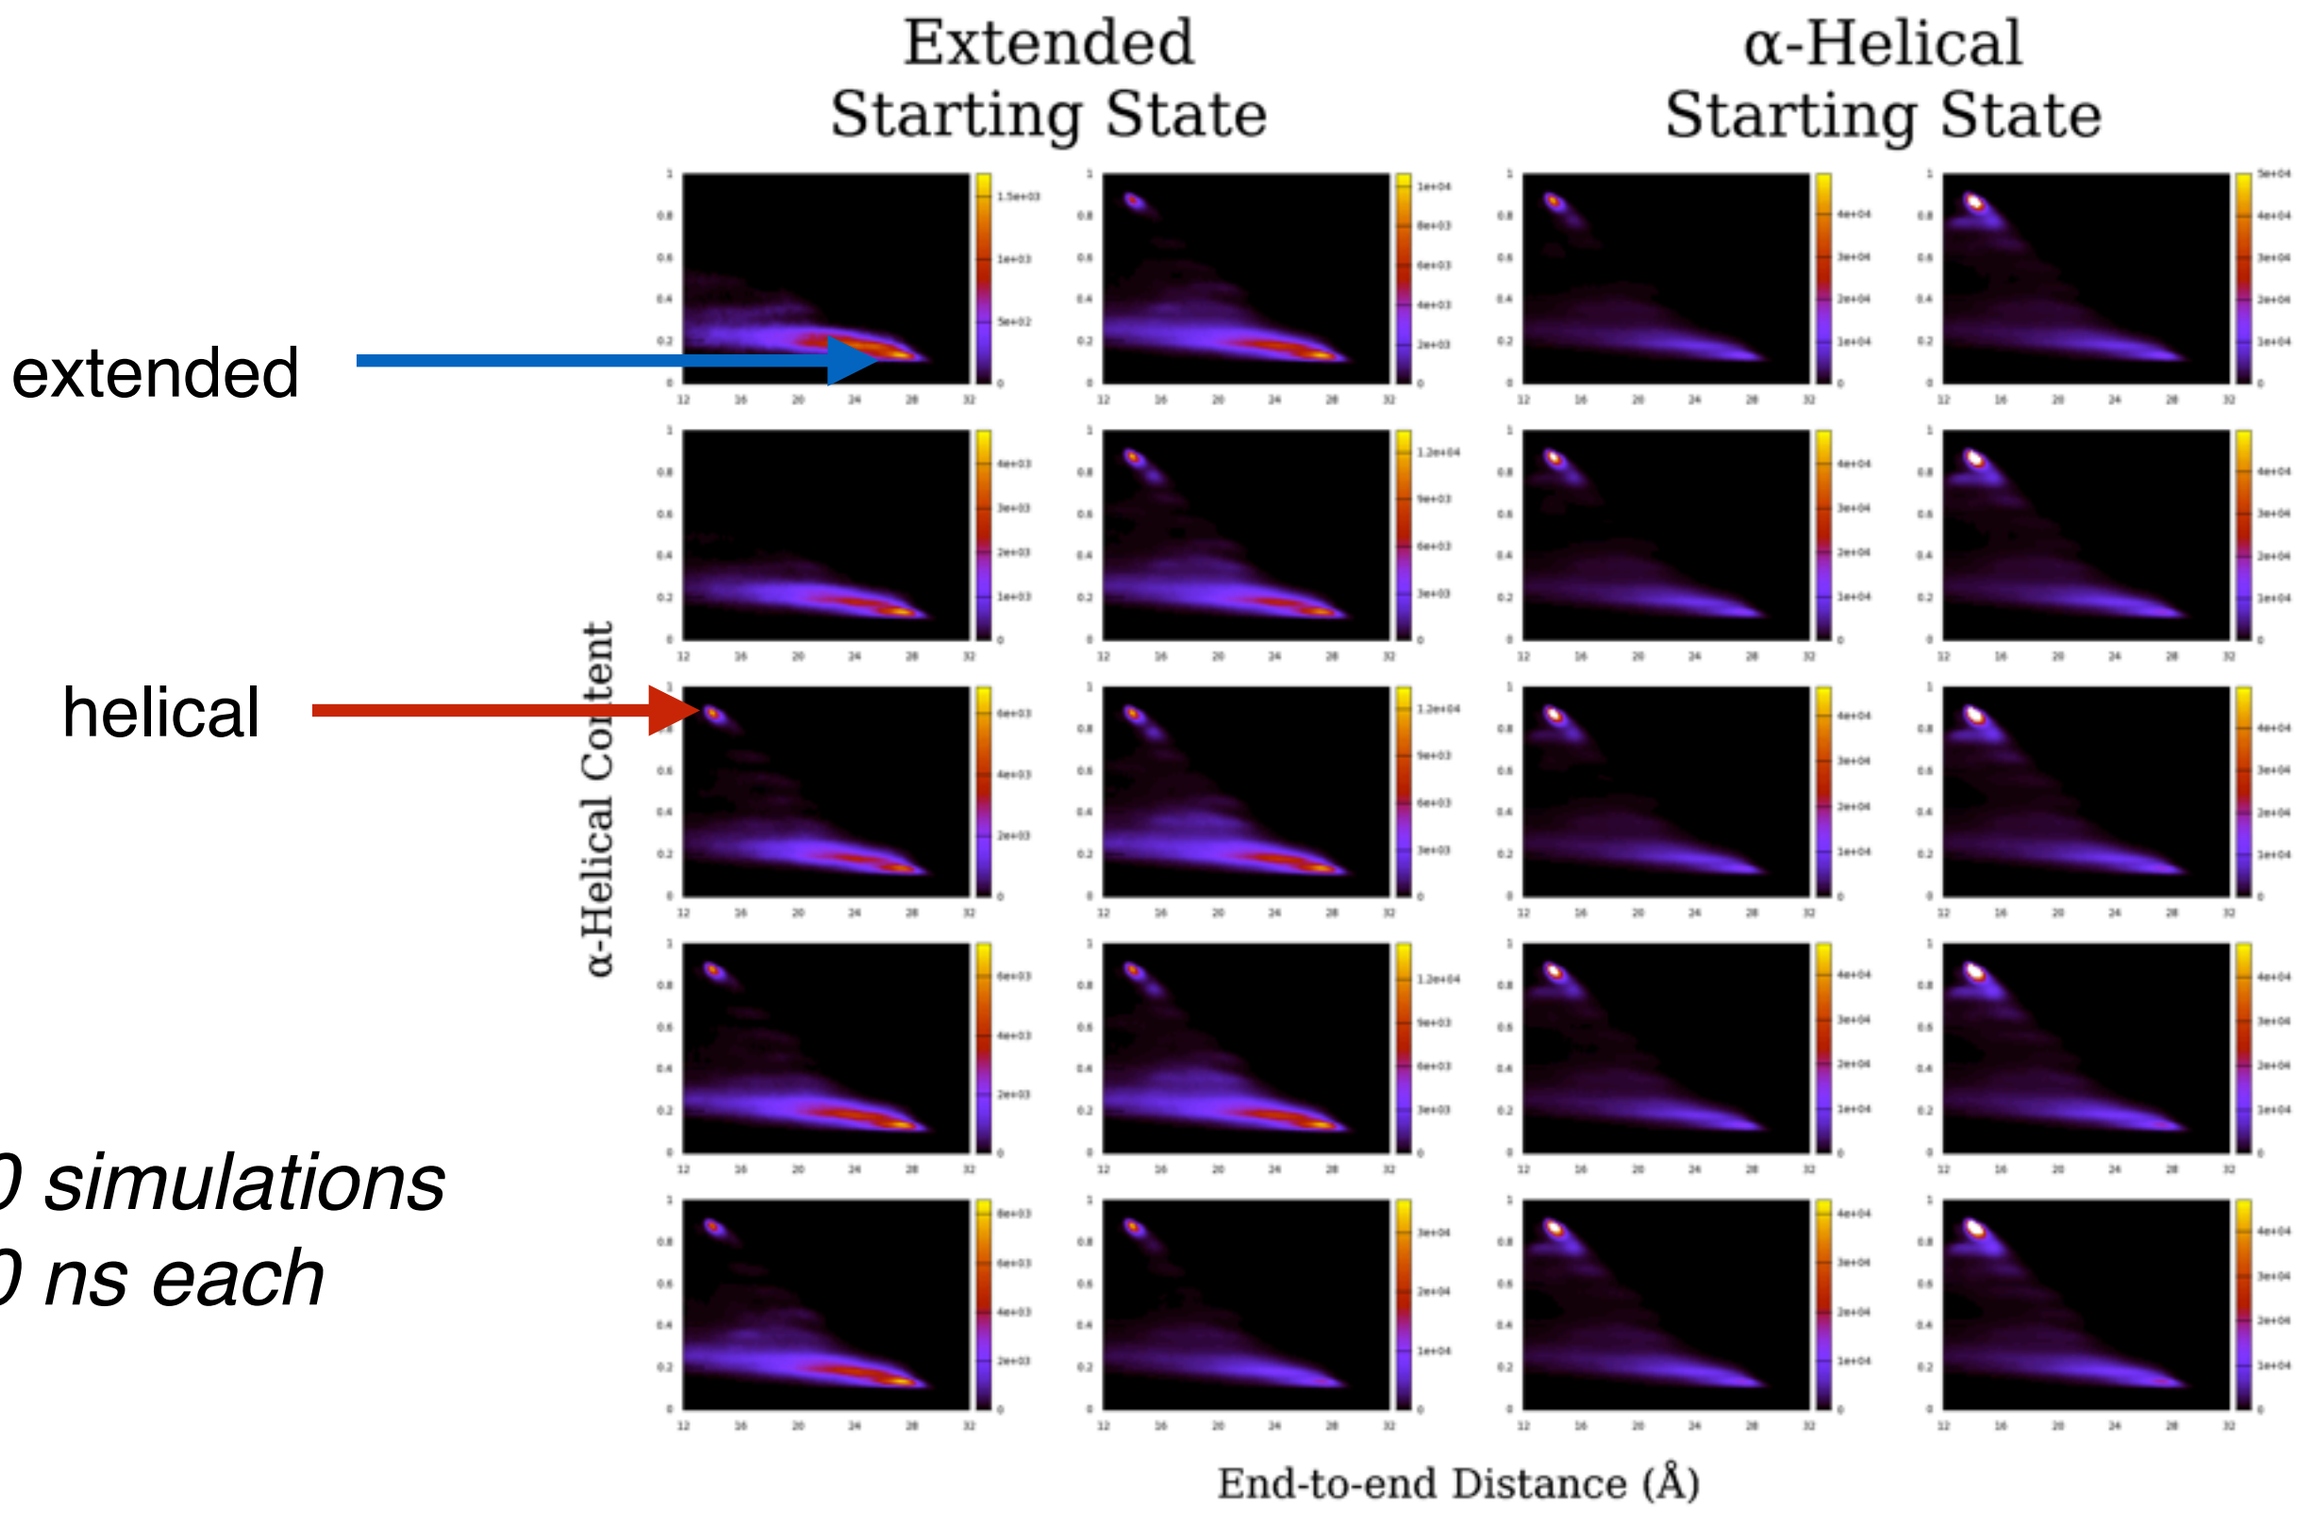

20 simulations
50 ns each

in almost every run, the peptide samples **both** helical and extended states

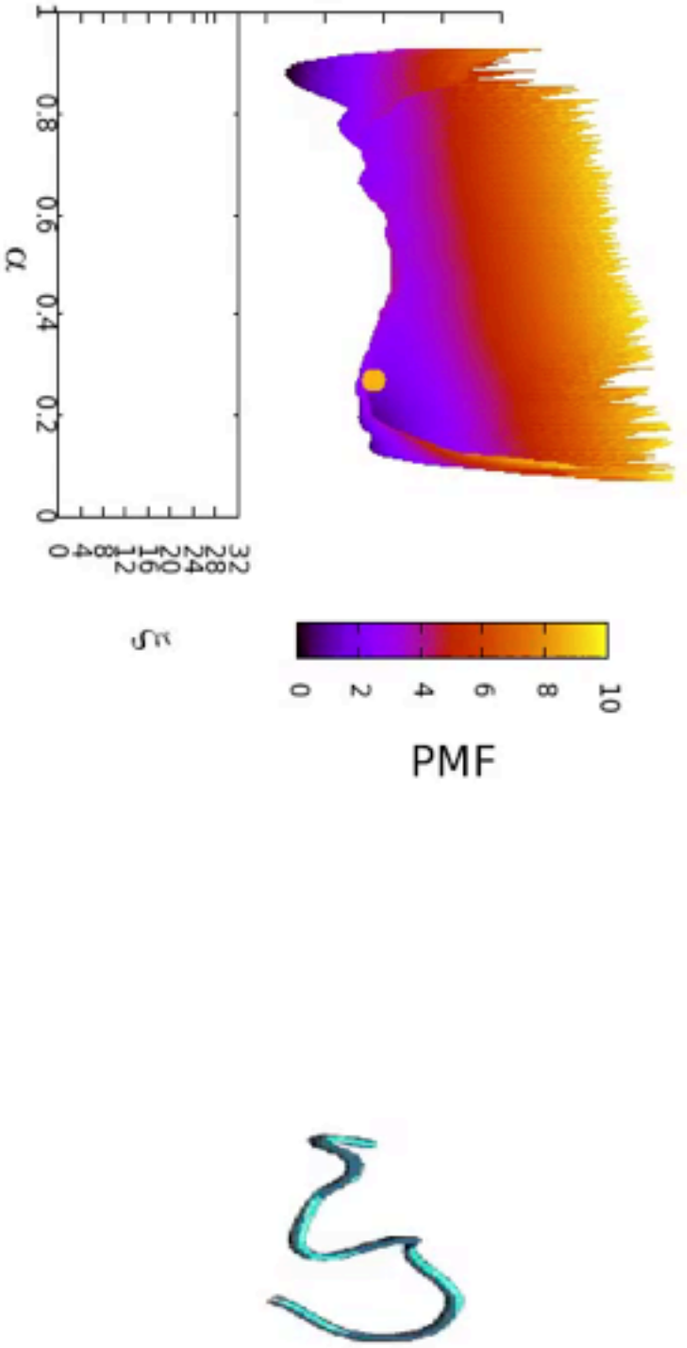
PMF
PMF

Example 2: OM protein development

insertion via BamA

hypothesis: folding and insertion are coupled via a strand-strand complement mechanism

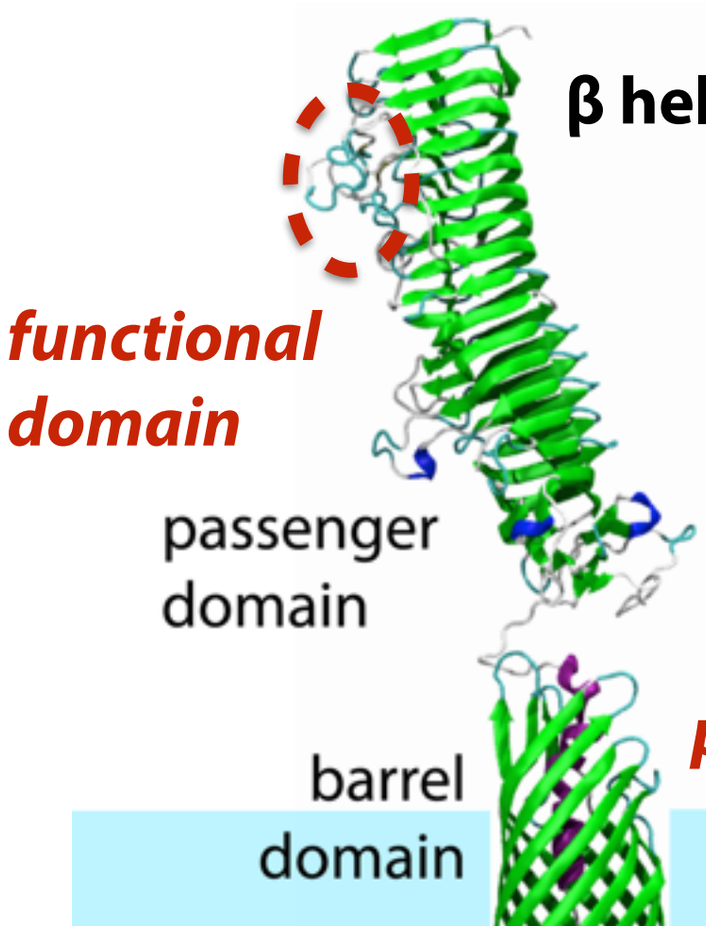

97% of autotransporters have a β helical domain that is **NOT** related to their virulence function - **WHY???**

hypothesis: folding drives export

simplest model for outer-membrane protein folding

GB1 - a β hairpin

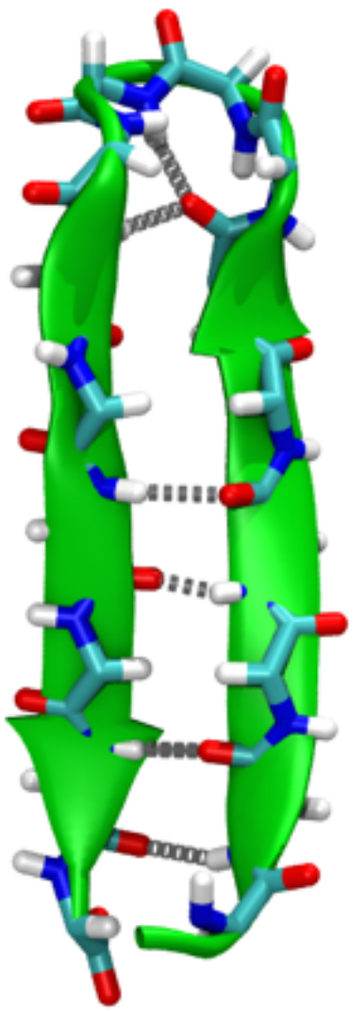

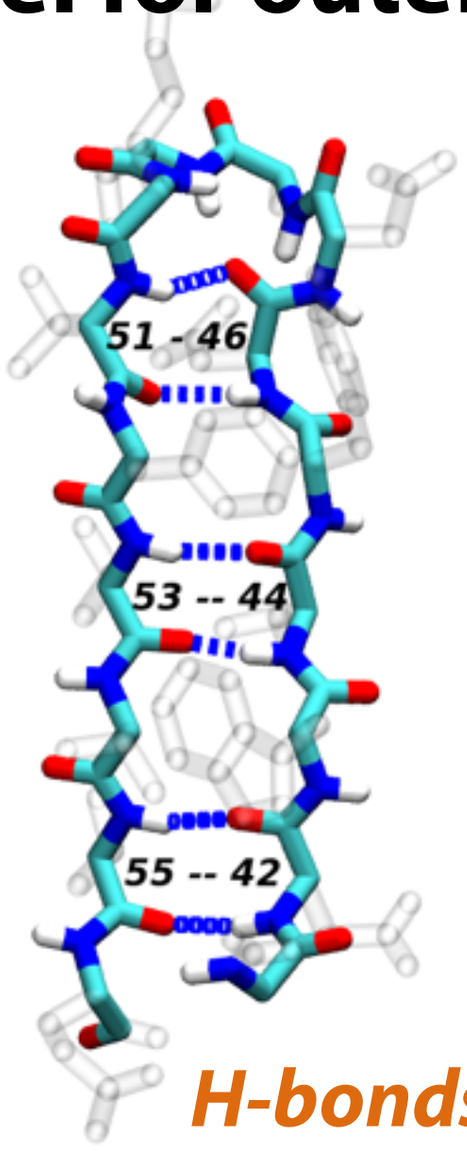

H-bonds

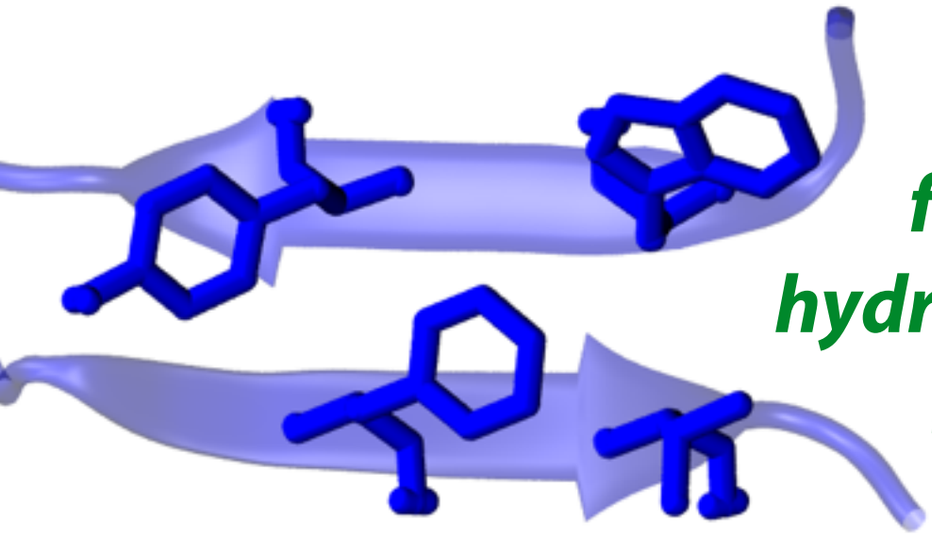

four core hydrophobic residues

Again, use two reaction coordinates

- R_G for **hydrophobic core** (4 residues), a measure of size
- backbone **hydrogen bonds** (up to 6), a measure of structure

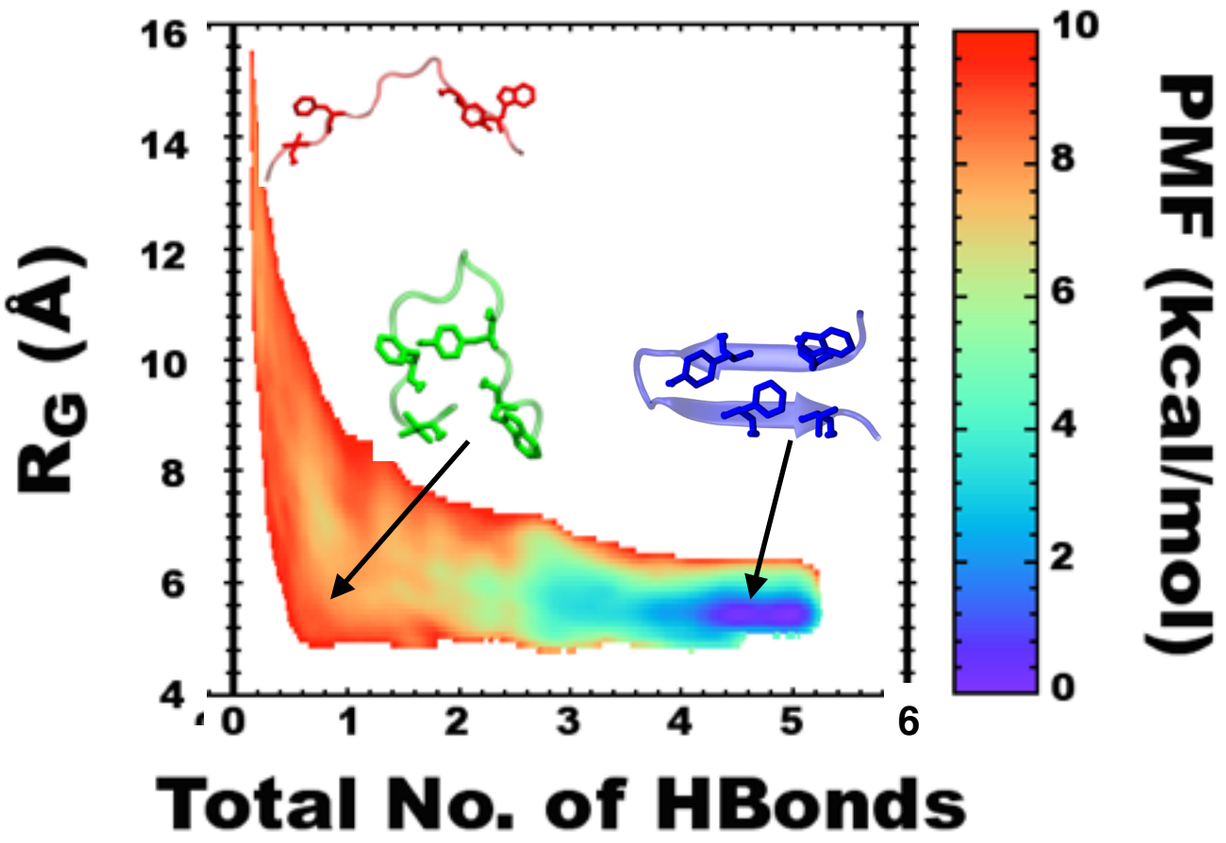

C36 favors hairpin by **8 kcal/mol!**

Comparison across force fields

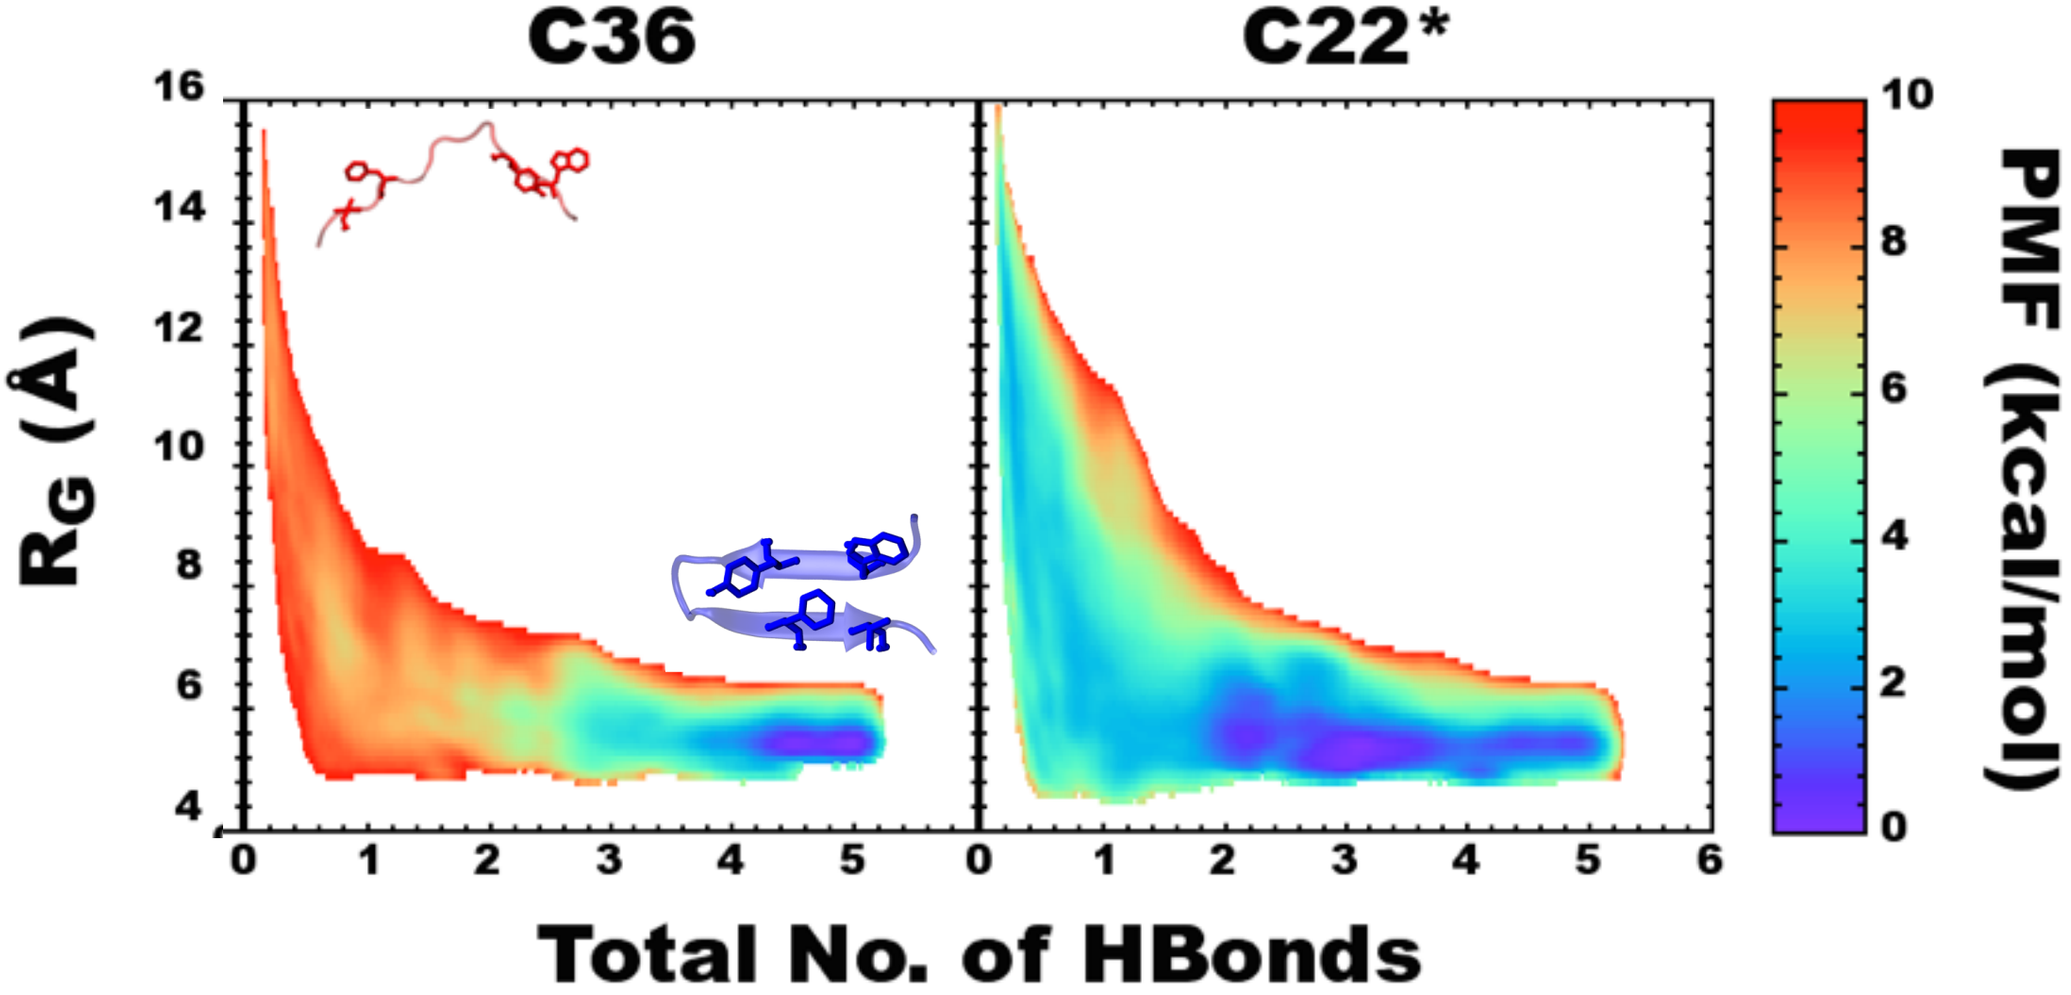

C22* favors hairpin by ~5 kcal/mol, a noticeable reduction

CHARMM36 known to **overstabilize** folded state (folded fraction of 77% vs. 60% experiment)

Best,…,MacKerell. *JCTC*. **8:**3257–3273**.** 2012.

Here, CHARMM22* may be best choice for observing folding transitions

What about Drude polarizable force field?

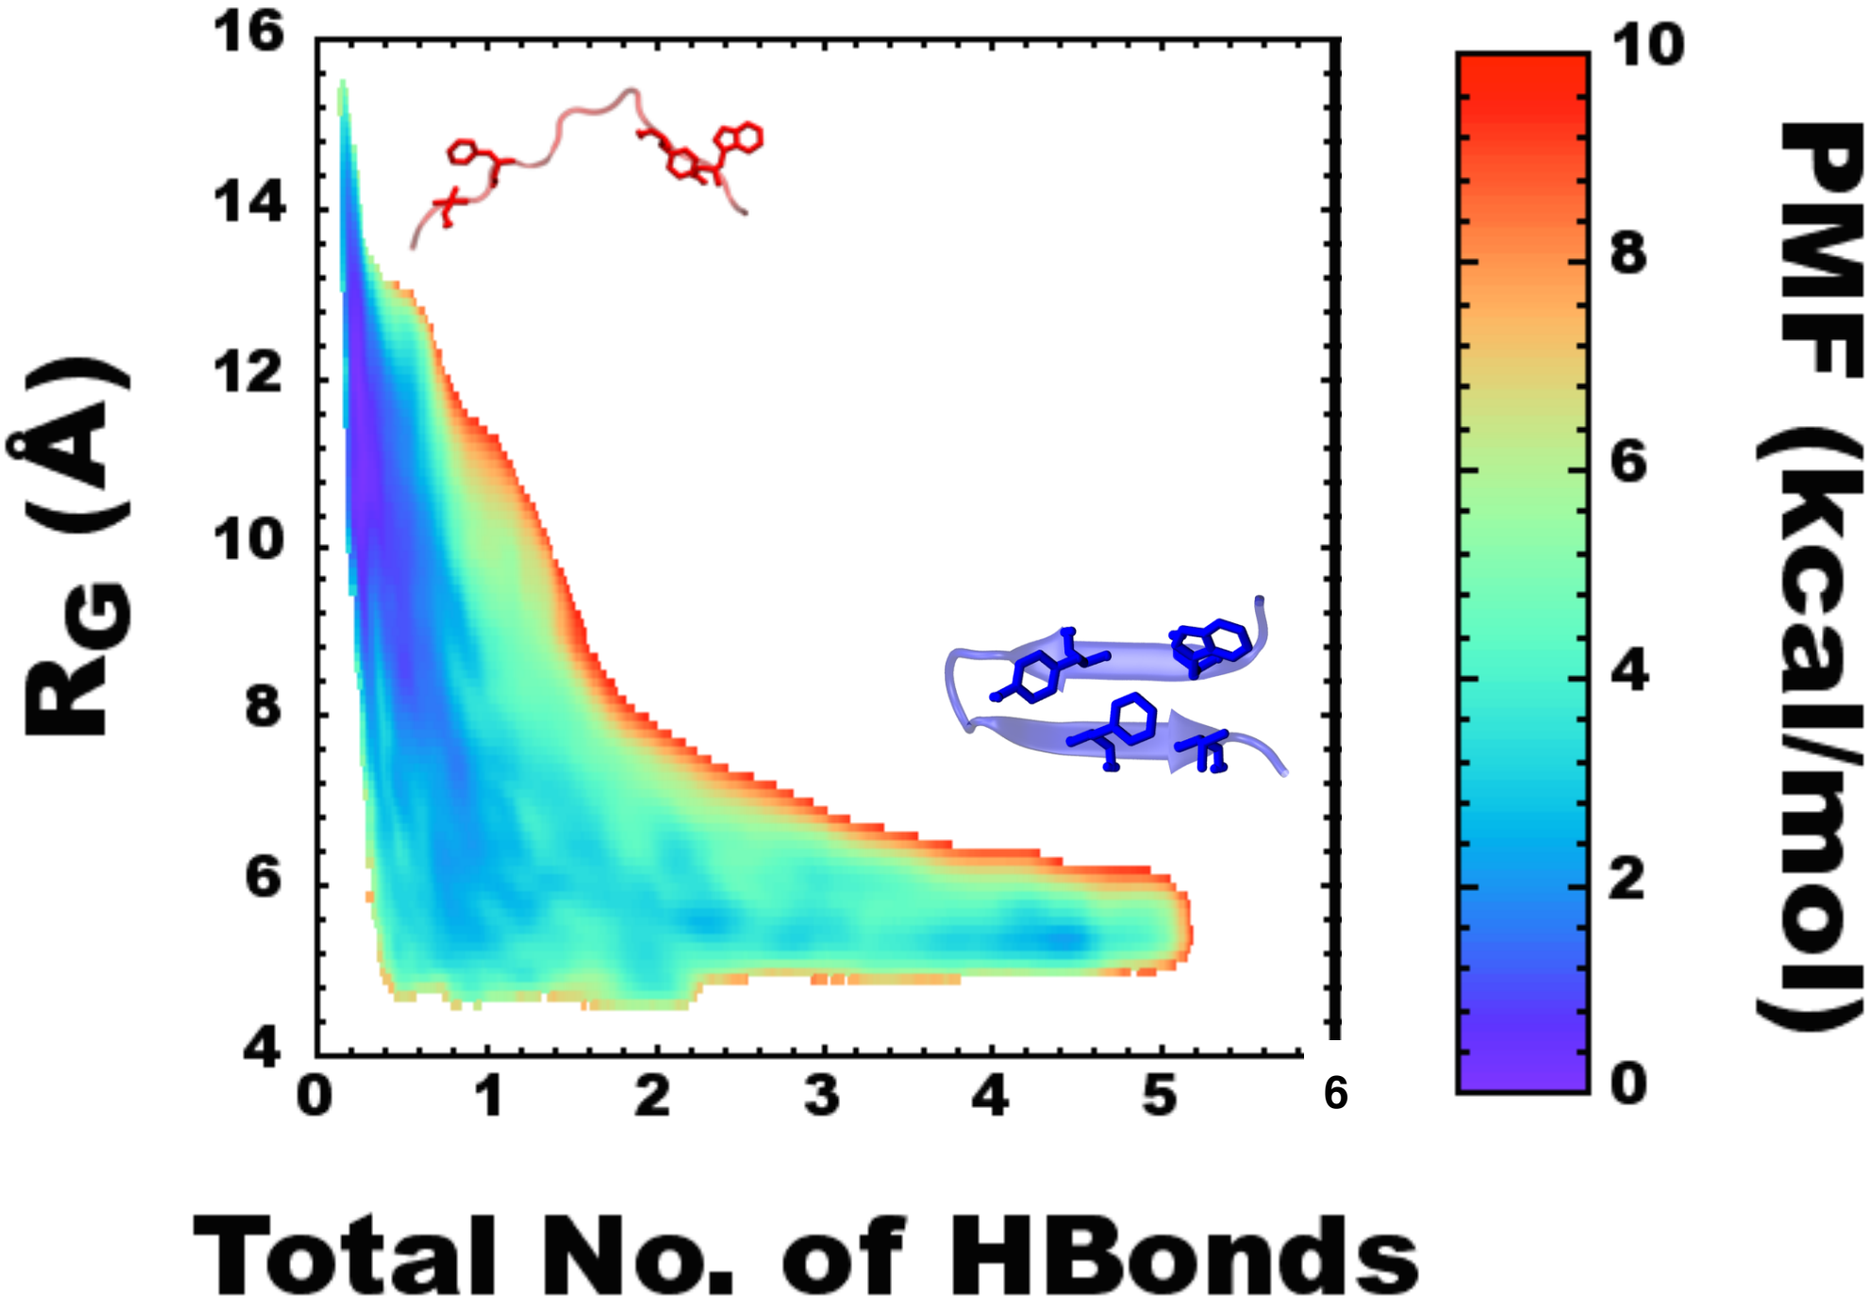

A **disaster**! Unfolded state is favored by 5-6 kcal/mol!

What is the source?

What about Drude?

pathway seems to be the same - cannot appeal to vastly different underlying states

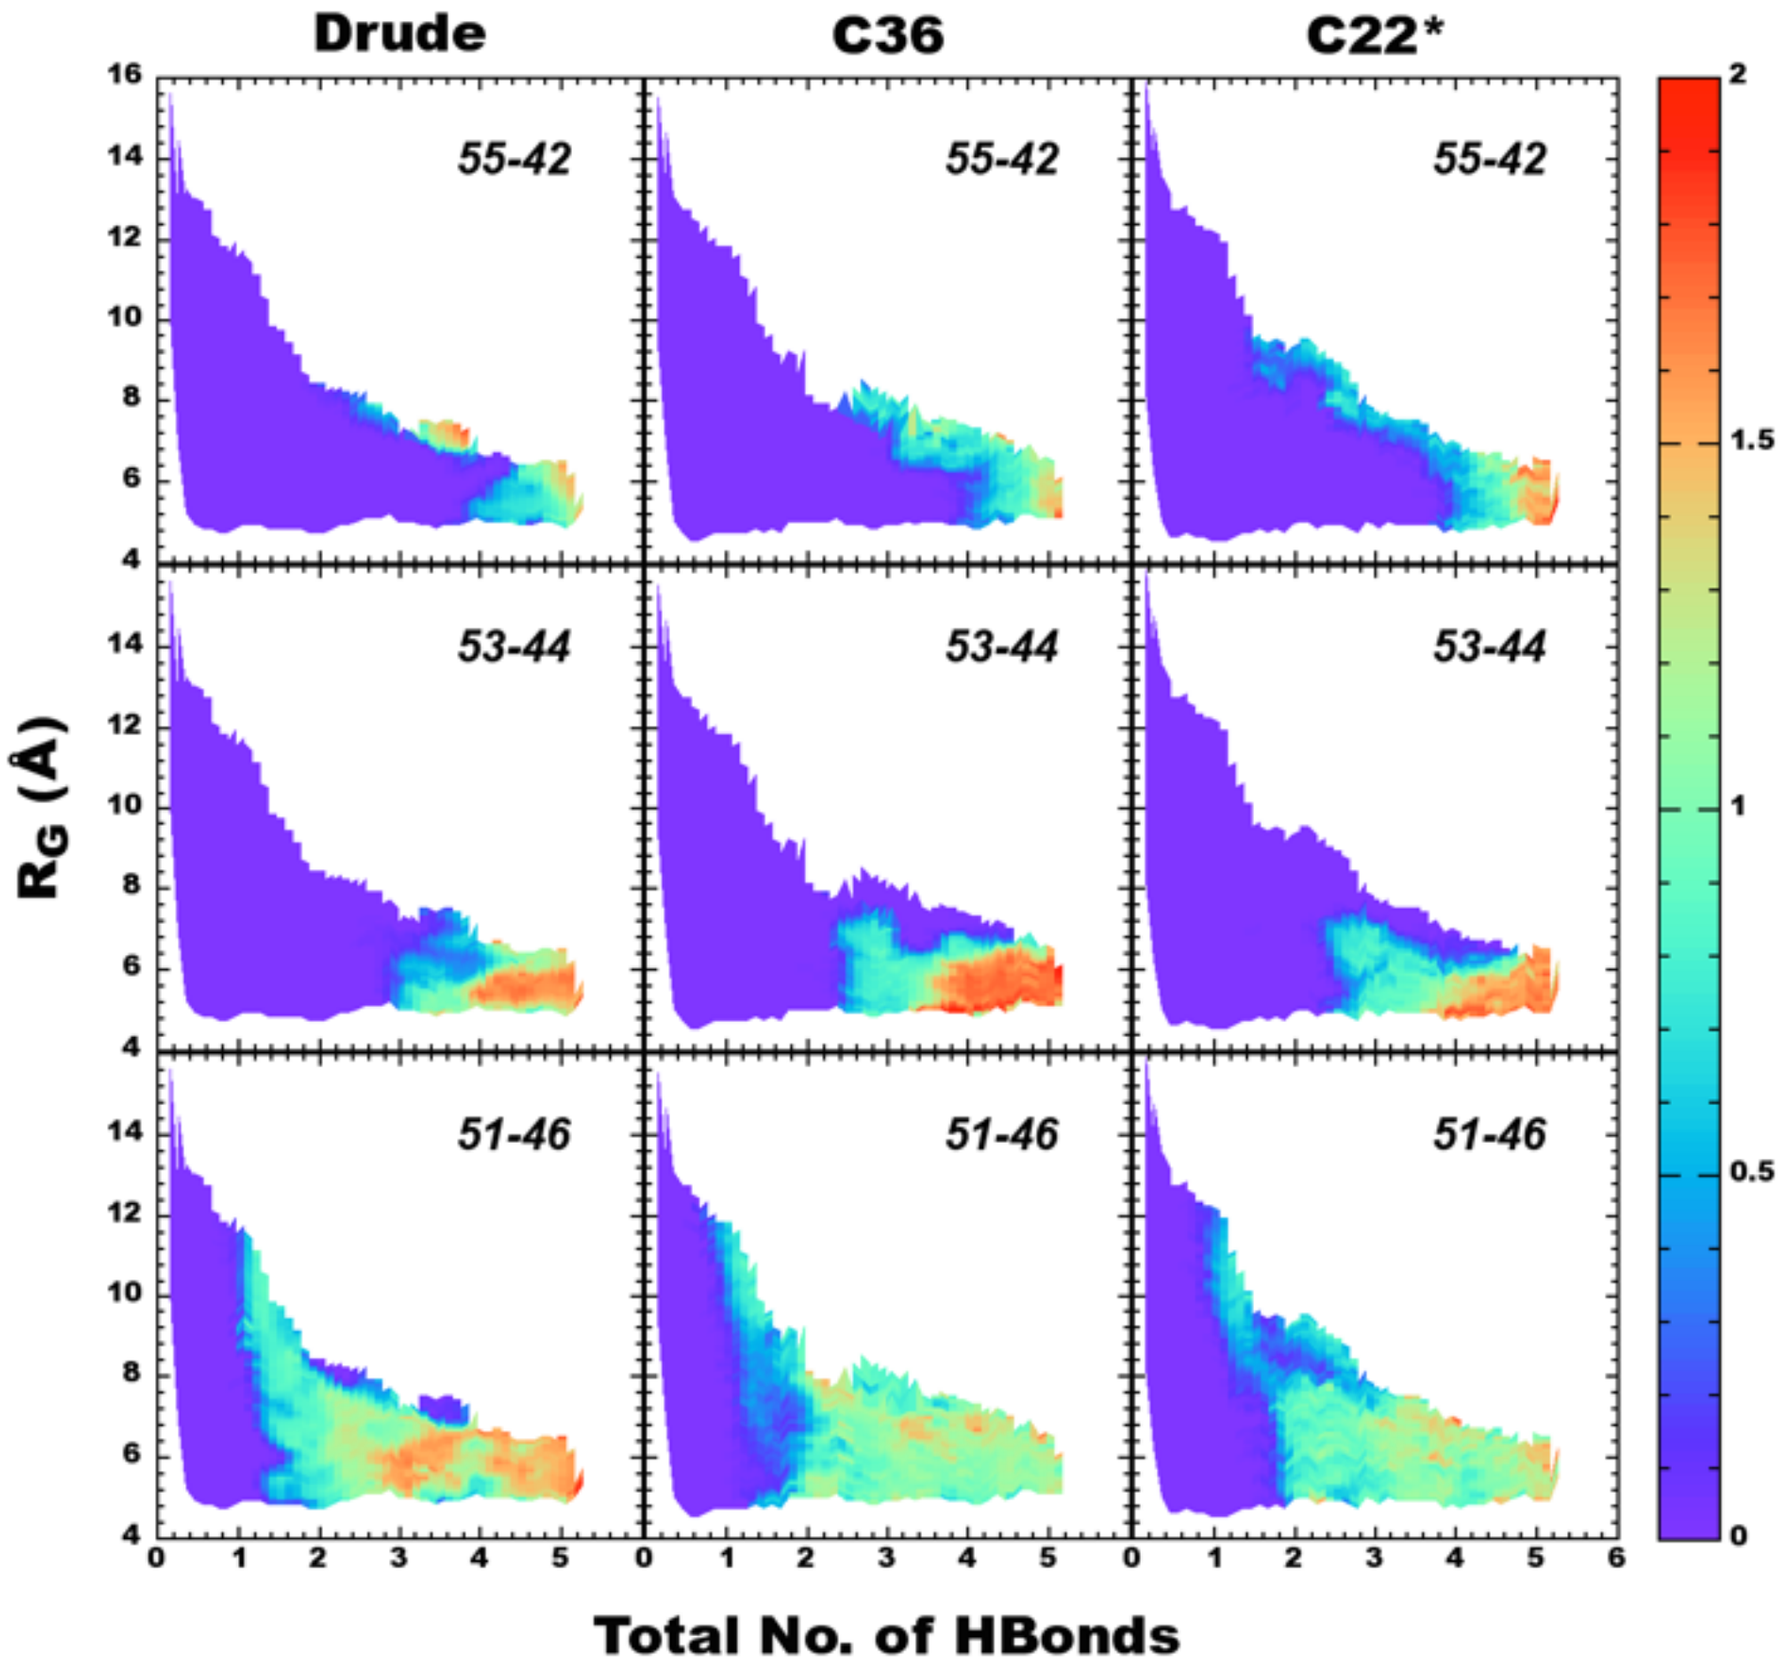

Energy breakdown

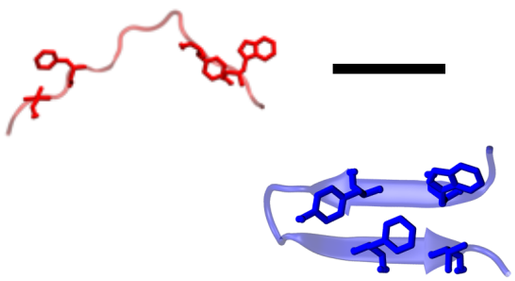

	Electrostatic	
	Drude	C36
$\Delta H_{\text{Extended}-\text{Hairpin}}$		
β-β	+22.2841	+13.6653
β-Water	–34.1321	–27.9719
β-Ions	+19.1761	+19.9284
Core-Core	+0.1122	–0.1911
Core-Water	–10.6076	+0.8959
Core-Ions	+1.1662	+0.8232
Sum	–2.0011	+7.1498
$\Delta H_{\text{Extended}-\text{Compact}}$		
β-β	–1.4706	–1.3416
β-Water	+56.1550	+56.3156
β-Ions	+3.8539	+4.2264
Core-Core	–1.0638	–0.8021
Core-Water	–40.6795	–27.5980
Core-Ions	+1.5464	+1.1326
Sum	+18.3414	+31.9329
$\Delta H_{\text{Compact}-\text{Hairpin}}$		
β-β	+23.7547	+15.0069
β-Water	–90.2871	–84.2875
β-Ions	+15.3222	+15.7020
Core-Core	+1.1760	+0.6110
Core-Water	+30.0719	+28.4939
Core-Ions	–0.3802	–0.3094
Sum	–20.3425	–24.7831

hydrophobic core residues' interactions with water are **too strong** in the extended state

are they becoming over-polarized?

ongoing work…

Example 3: membrane permeation

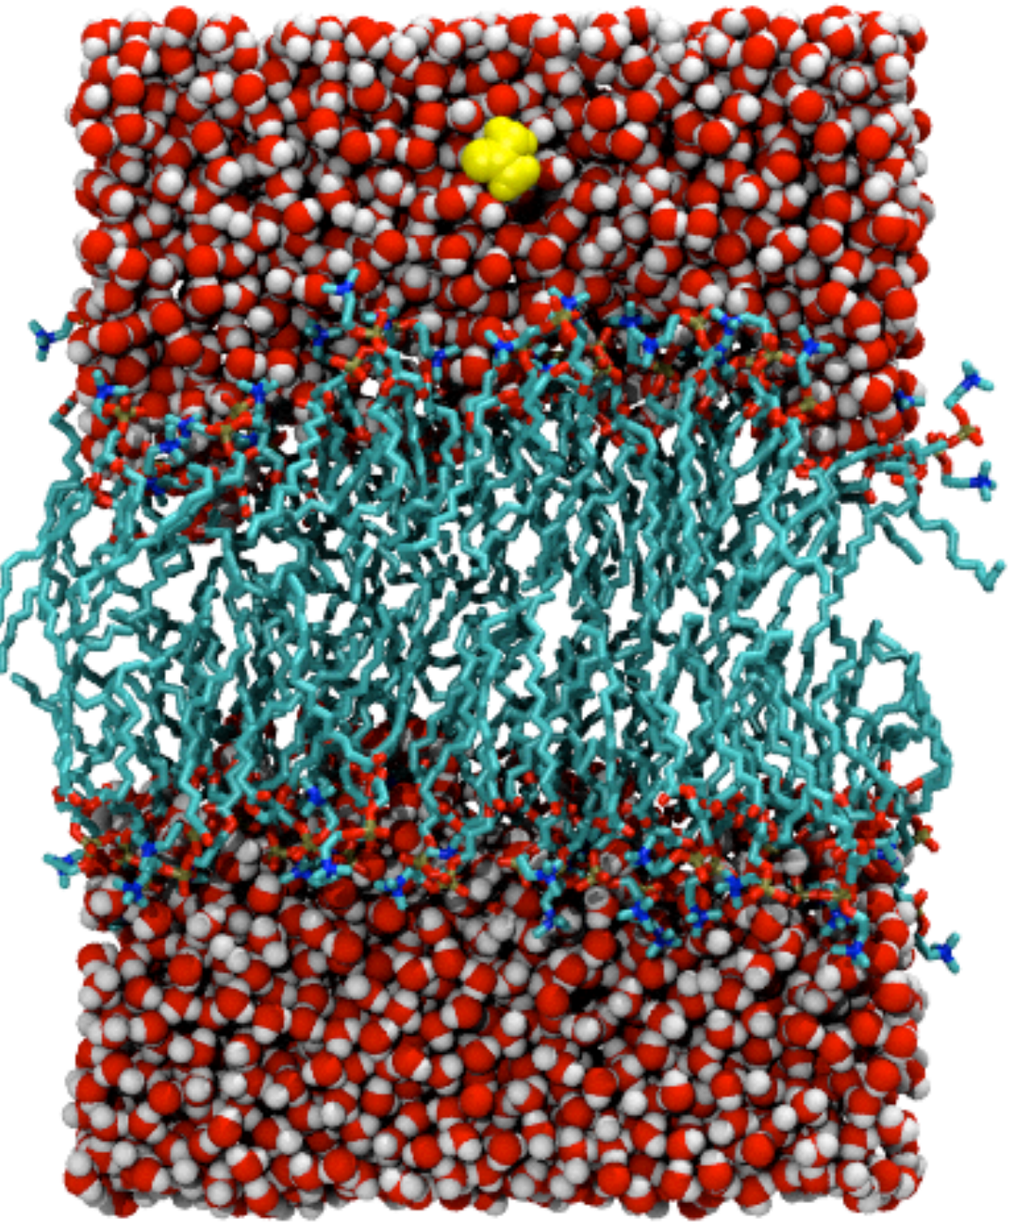

small molecules have a probability of breaching the membrane barrier related to their **potential of mean force W(z)** and **diffusivity D(z)** by the solubility-diffusion model:

$$\frac{1}{P} = \int_{\text{bulk}}^{z} \frac{e^{W(z')/kT}}{D(z')} dz'$$

note: focus is on the PMF only here, not the permeability

using urea, etc. as simple test cases for this model and its constituent calculations

**collaboration with Chris Chipot, Yi Wang, Chris Rowley, and Rommie Amaro*

Lee et al., Gumbart. Simulation-based approaches for determining membrane permeability of small compounds. (2016) *J. Chemical Information and Modeling*, 56:721-733.

How much sampling is needed?

is 10 ns/window is enough? (360 ns total)

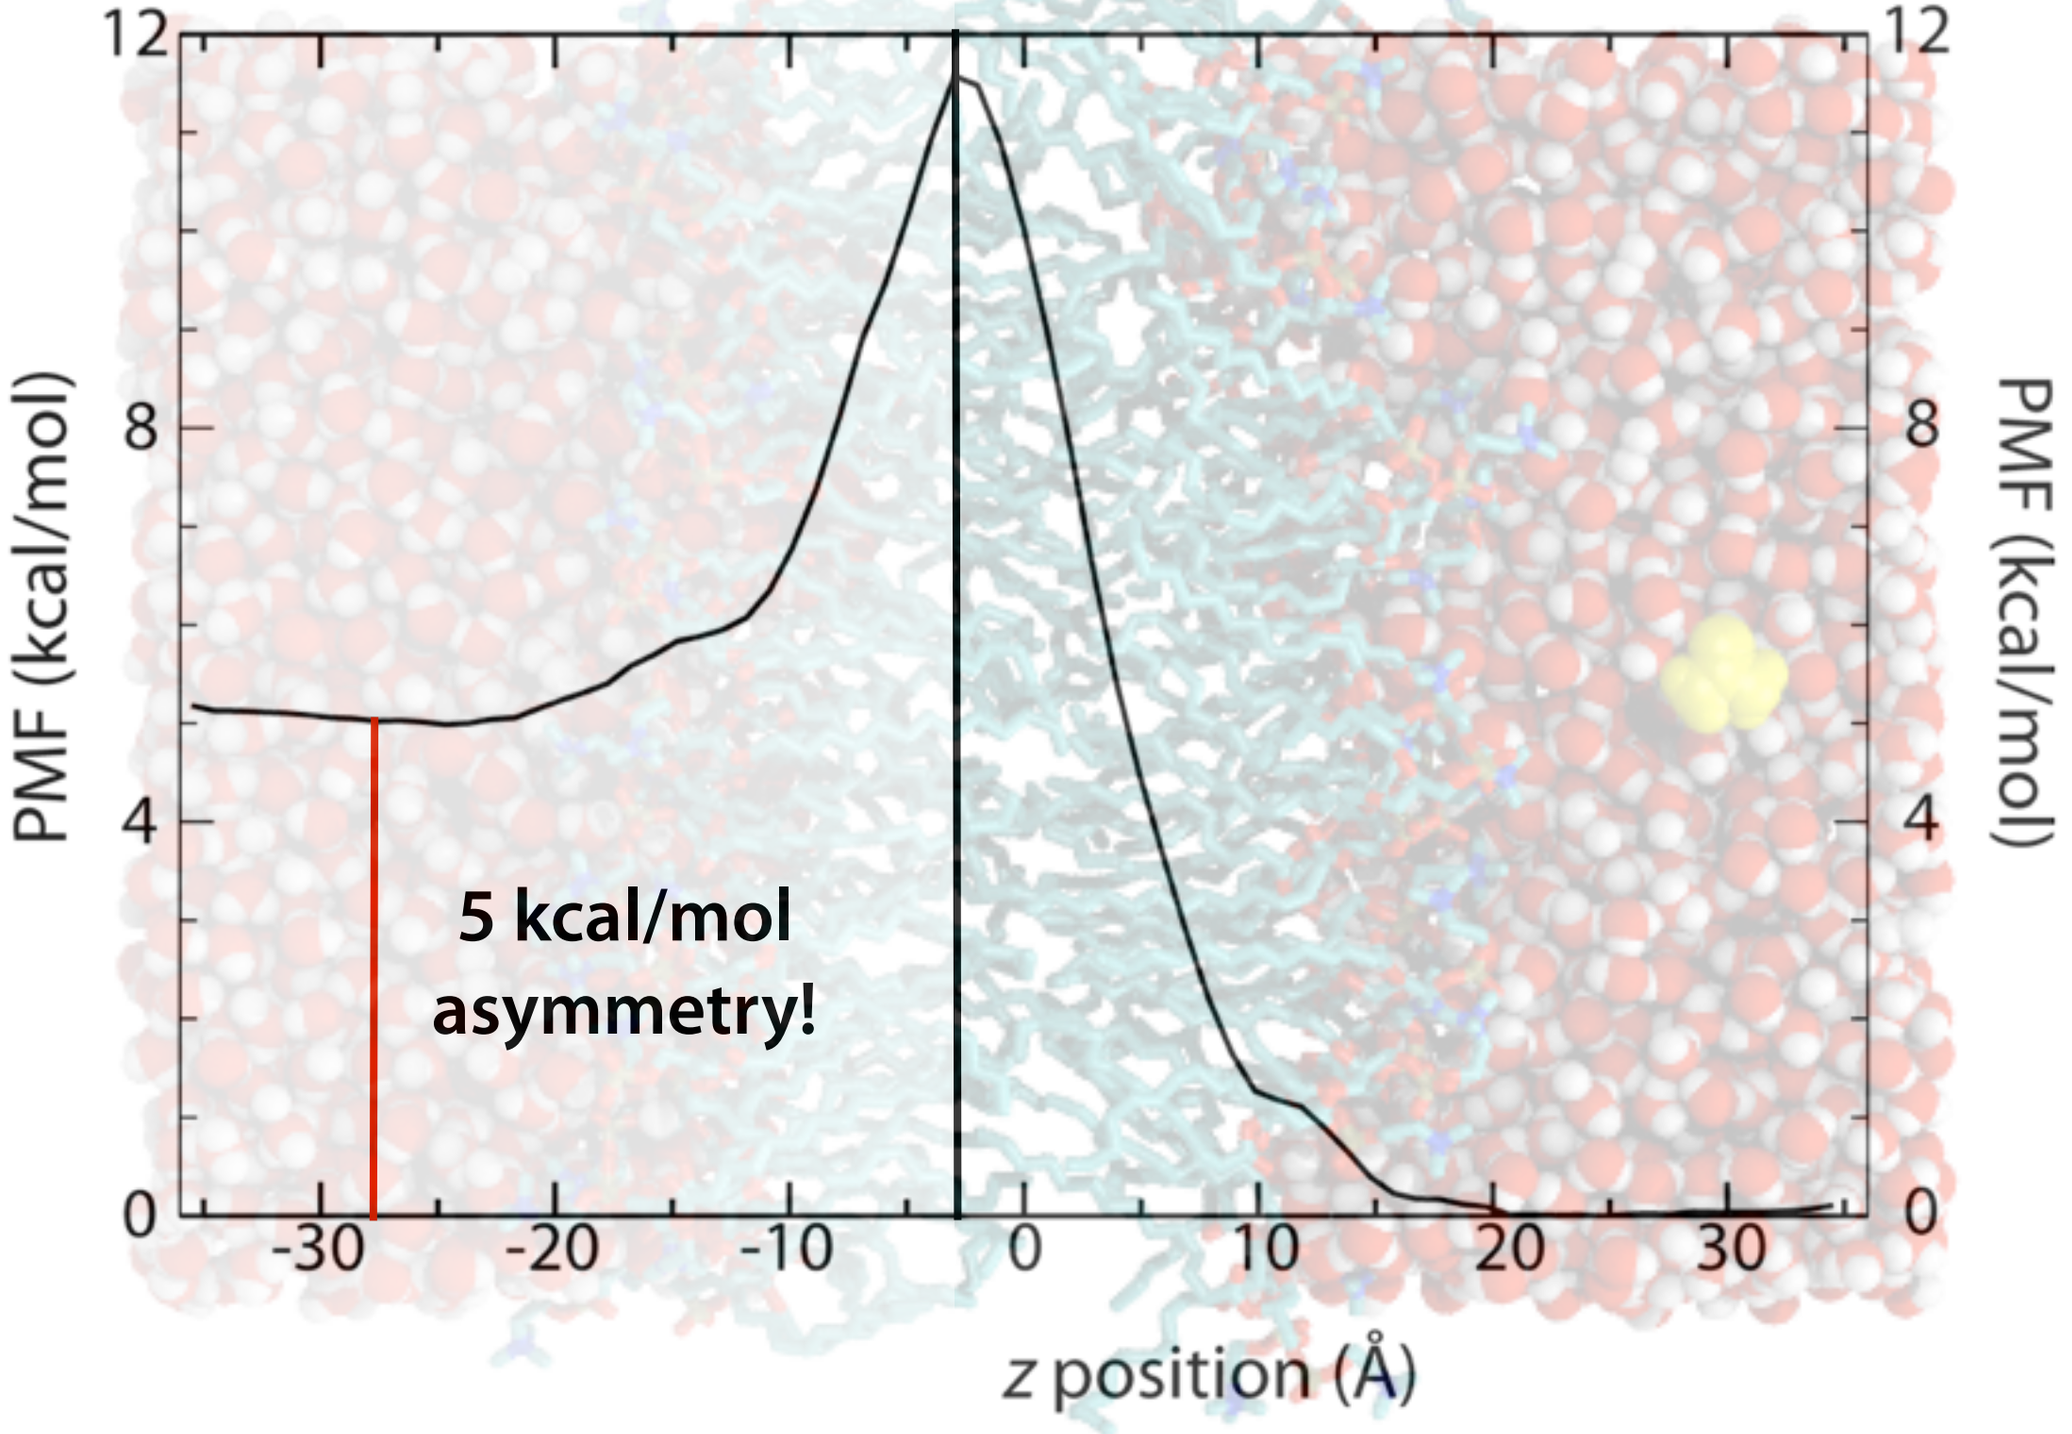

How much sampling is needed?

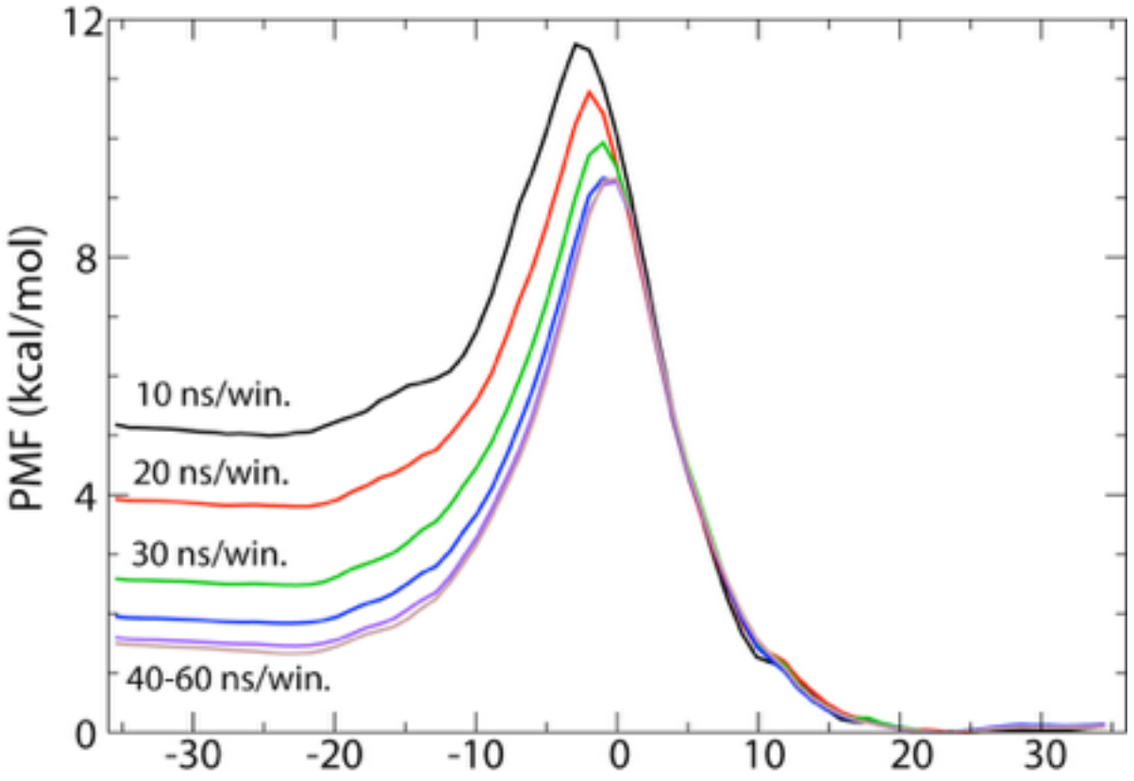

requires over **4 *μ*s** of REMD-US to bring asymmetry to ~1 kcal/mol but the results are good, **log(P) = -5.83**

exp. ~ **-5.4** (Finkelstein 1976), although from PAMPA, ~ -9.0 (not a membrane!)

ABF, on the other hand, requires less than 1 *μ*s for similar results - *why*?

ABF calculations run in the lab of Yi Wang, Chinese Univ. of Hong Kong

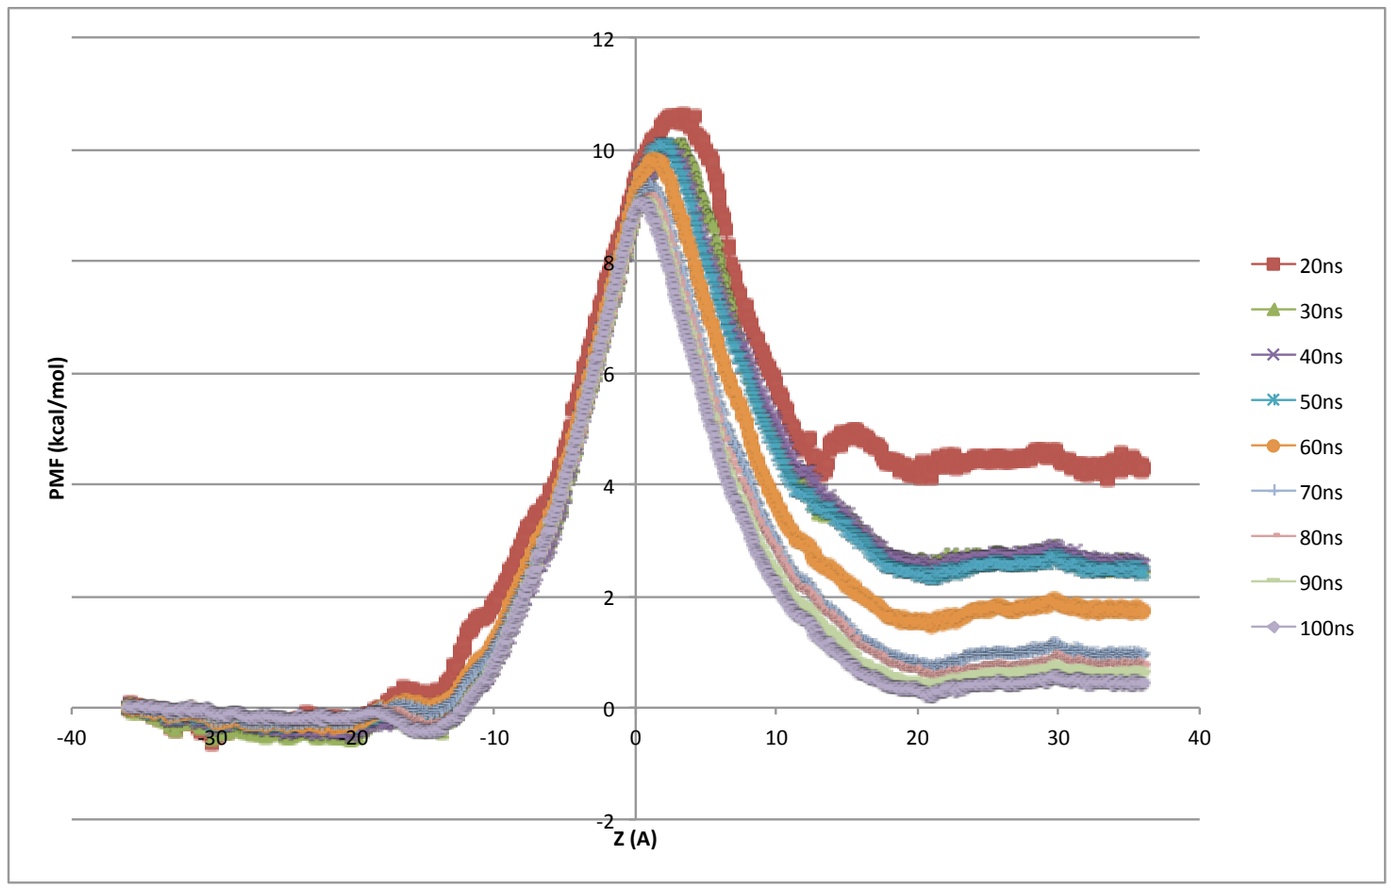

Can REMD-US be rescued here?

z = 25, 15 ns

z = 15, 50 ns

z = 5, 50 ns

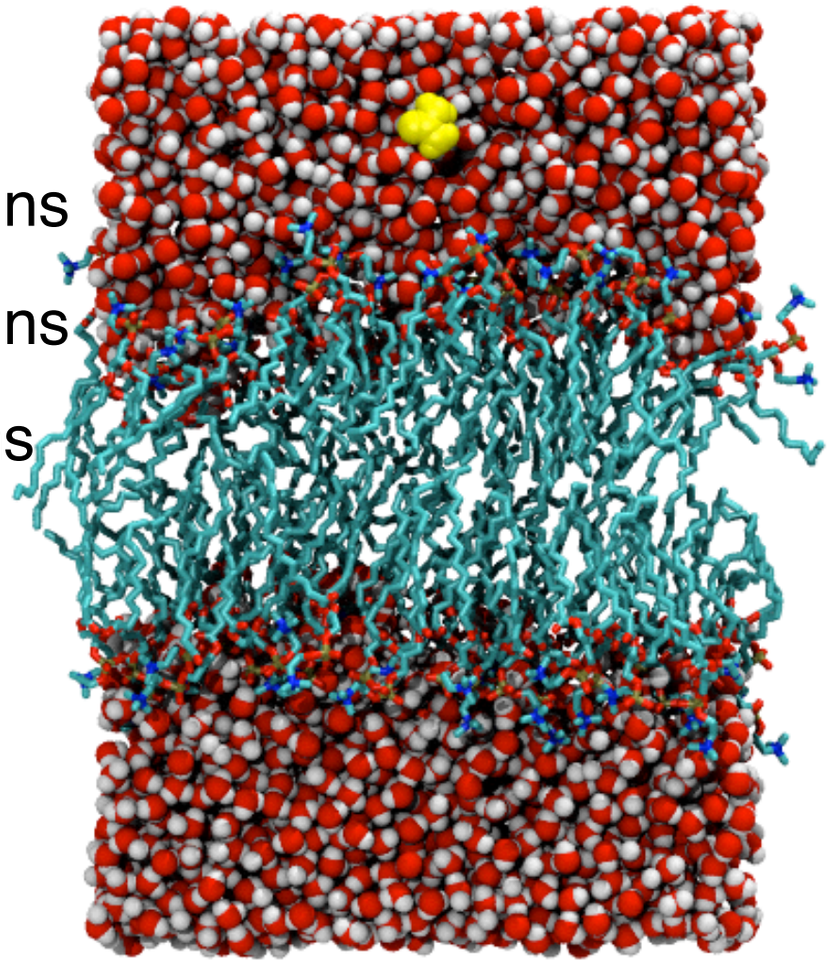

z = 30, 10 ns

z = 20, 25 ns

z = 10, 50 ns

z = 0, 100 ns

Equilibrate a few selected windows and use them to seed intermediate ones

Every 5 Å, both above and below, **500 ns** total

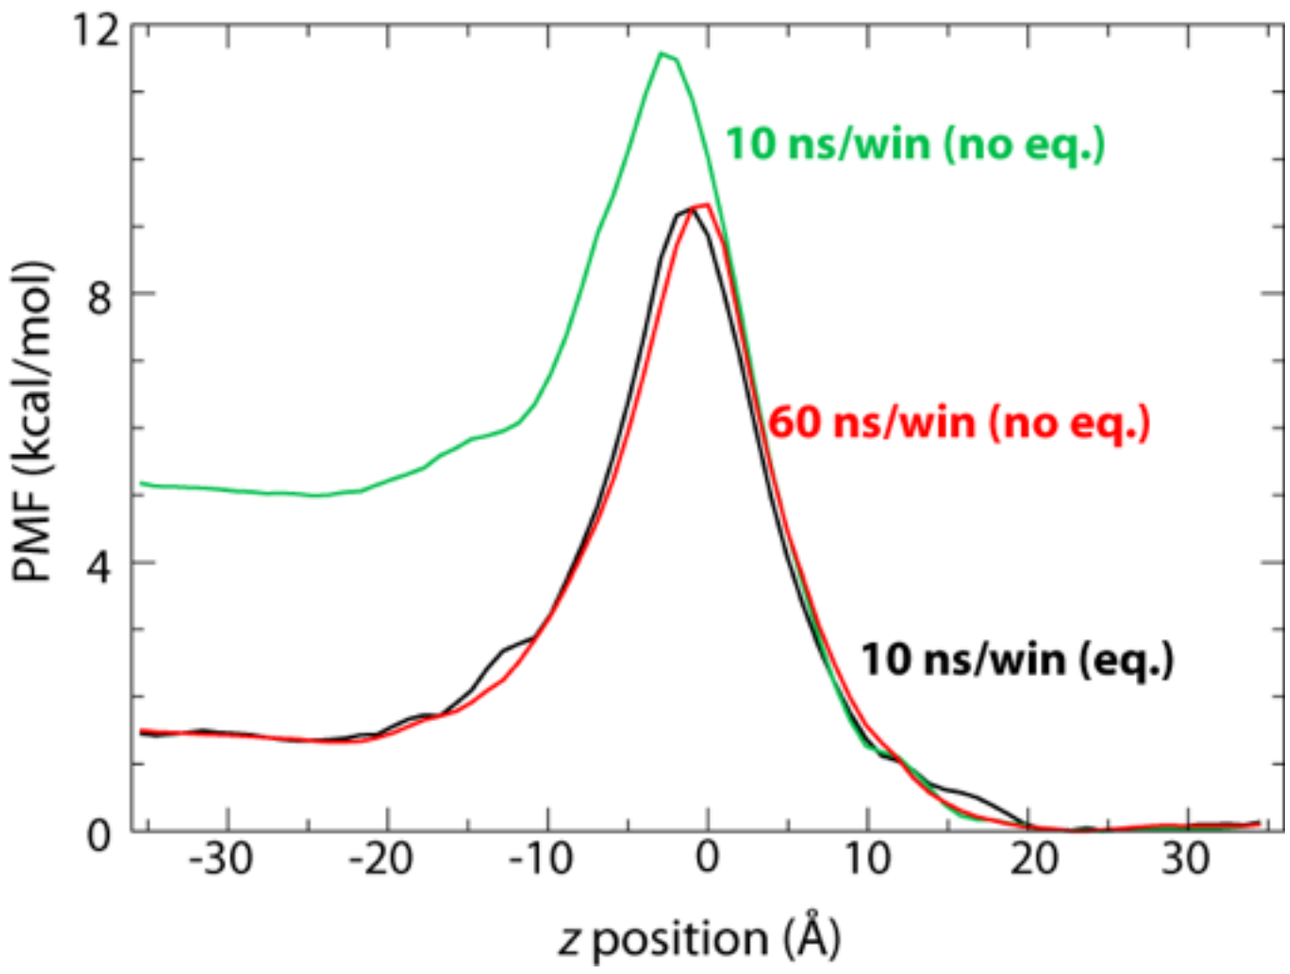

Using the newly **equilibrated** states, REMD-US produces an identical PMF in **~25%** of the time (500 ns + 720 ns = 1220 ns vs. 4320 ns)

Final thoughts and reminders

Biasing is fine as long as there is a method to **unbias** the results

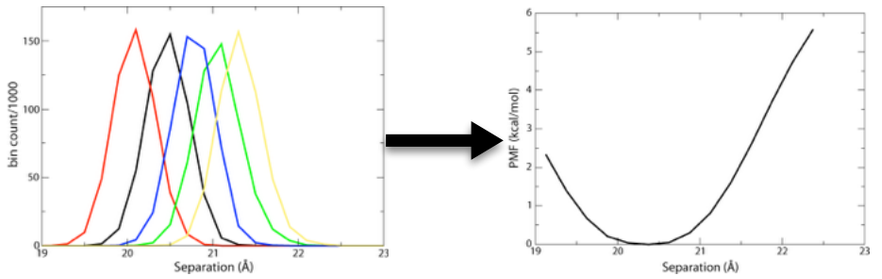

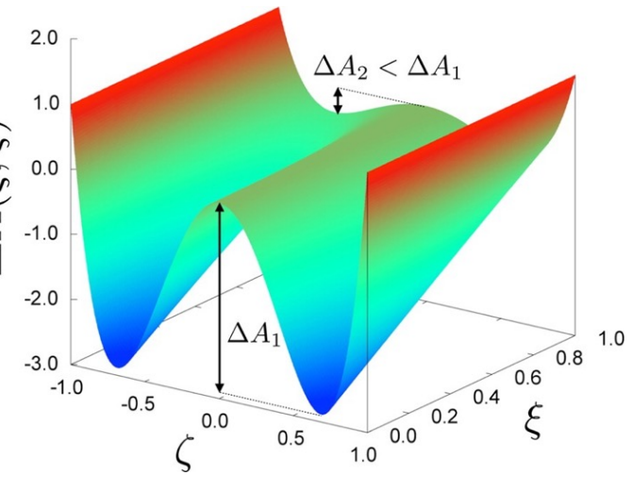

Be aware of **orthogonal barriers** that may hamper sampling in some regions

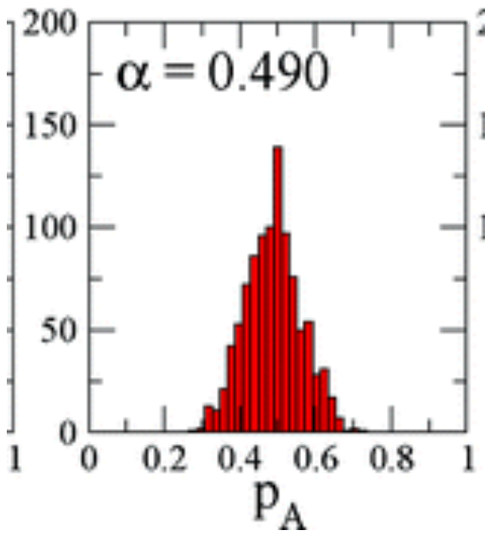

Choose **reaction coordinate(s)** that capture the transition of interest well and check it (e.g., with the **committor probability**)

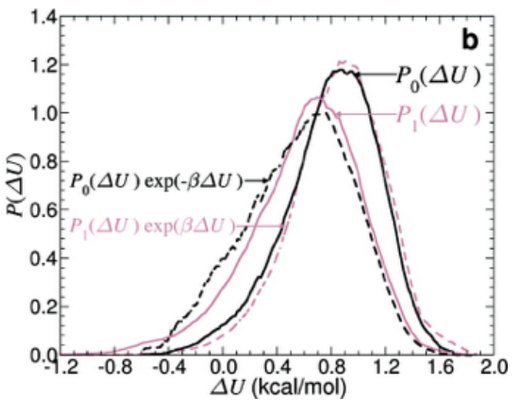

For any method involving stratification of a path (US, FEP, etc.) confirm that the neighboring states have good **overlap**

Be certain that your initial states are **well equilibrated**

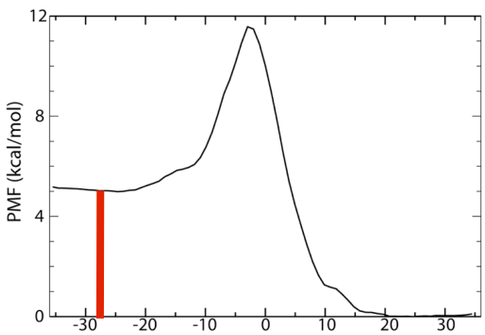

Free energy tutorials available

http://www.ks.uiuc.edu/Training/Tutorials

Free-energy perturbation

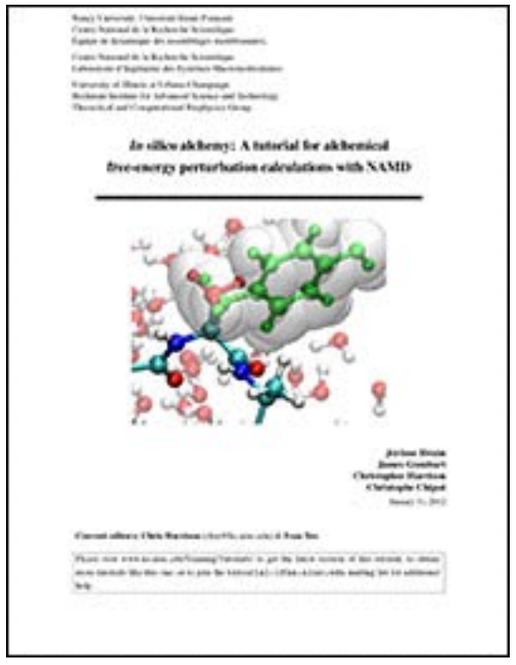

Protein:ligand binding

Umbrella Sampling, SMD

Adaptive Biasing Forces