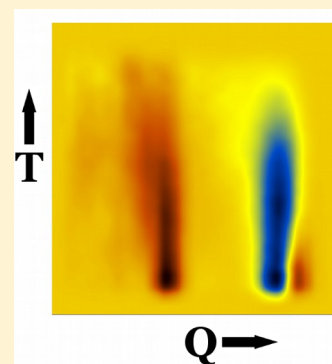# Sensitivity of Folding Molecular Dynamics Simulations to Even Minor Force Field Changes

Athanasia-Panagiota Serafeim, Georgios Salamanos, Kalliopi K. Patapati, and Nicholas M. Glykos*

Department of Molecular Biology and Genetics, Democritus University of Thrace, University campus, 68100 Alexandroupolis, Greece

S *Supporting Information*

ABSTRACT: We examine the sensitivity of folding molecular dynamics simulations on the choice between three variants of the same force field (the AMBER99SB force field and its ILDN, NMR-ILDN, and STAR-ILDN variants). Using two different peptide systems (a marginally stable helical peptide and a β-hairpin) and a grand total of more than 20 μs of simulation time we show that even relatively minor force field changes can lead to appreciable differences in the peptide folding behavior.

1. INTRODUCTION

Folding molecular dynamics simulations of peptides and proteins are an acid test for the ability of biomolecular empirical force fields to reproduce the physical reality.[1−12] The reason is that in the case of a folding simulation the whole complexity of the peptide's folding landscape must be reproduced and not just the dynamics of the biomolecule's folded state. What this line of thought implies, then, is that folding simulations can be used as a sensitive tool to gauge not only the ability of any given force field to reproduce the experimental data but also as a purely comparative tool to quantify the sensitivity of folding trajectories to even minor force field changes. We believe that assessing the sensitivity of molecular dynamics trajectories to relatively small force field variation is important for quantifying the convergence − or otherwise − of force fields: if within the same force field family significant differences in the folding behavior are observed, then this would demonstrate the complexity and delicate balance required to fine-tune biomolecular force fields and would serve as a reminder of just how difficult it is to develop a physically relevant biomolecular force field.

In this communication we use three variants of one of the most successful and popular biomolecular force fields − the AMBER99SB force field[13,14] − to quantify the sensitivity of the folding trajectories on relatively small force field changes. The three variants that we selected to study (namely, the ILDN, NMR-ILDN, and STAR-ILDN variants, see next section) essentially differ only in the backbone dihedral angle potential with these differences being due and arising from the different strategies used for their respective optimization (see references in the next section).

Given the very large number of recent papers that compare the performance of different force fields (see, for example refs

6, 7, and 10 and references therein) a natural question arises as to the significance of the present contribution. We believe that the present study is a useful addition to the existing literature for several reasons. The first is that we have used two peptide systems that differ significantly in terms of the secondary structure of their respective folded states (helical vs β-hairpin). The second is that the peptides studied here also differ markedly in the stabilities of their folded states (the β-hairpin is the prototype of a fast folding and stable peptide, the helical peptide is marginally stable). The third reason is that by choosing such a closely related set of force fields, any differences observed in the peptides' folding behavior can immediately be attributed to their variation (which essentially is the backbone dihedral angle potential). The fourth reason is that we have selected peptides that are already known to fold to the correct secondary structure using the parental force field and can, thus, minimize the complications arising from using a desperate set of force fields. Lastly, by selecting relatively small peptides and performing a total of 20 μs of adaptive tempering simulation we can possibly establish statistical significance for the observed differences. In the paragraphs that follow we discuss the design, implementation, and statistical significance of the calculations performed, the results obtained from the simulations, and the implication of these results for force field convergence and development.

2. METHODS

2.1. Peptides and Design of the Simulations. We have chosen to simulate two well-studied peptides that are already known to fold correctly using the plain AMBER99SB force

field. The first peptide is a variant of residues 101−111 of the human α-Lactalbumin protein (peptide sequence INYWLA-HAKAG). This peptide, hereafter denoted as the αLa peptide, is known from two independent NMR determinations[15,16] to possess a marginally stable 3_{10}-helical structure in its N-terminus (residues 3−6) with its C-terminus being mostly disordered and is known to fold to a helical structure using the plain AMBER99SB force field.[17] The second peptide is CLN025 which is the prototype for a fast folding and very stable β-hairpin.[18] The difference in the stabilities of the two peptides necessitated the application of two different simulation protocols: Whereas αLa − and due to its marginal stability − could be studied with a classical NpT simulation, for the very stable CLN025 peptide we resorted to adaptive tempering[19] as a means to increase the number of folding/unfolding events and to improve the sampling of the non-native states. The three force field variants that we used for our simulations are (a) the AMBER99SB-ILDN force field,[20] hereafter referred to as 'ILDN', (b) the AMBER99SB-NMR-ILDN variant,[21] hereafter referred to as 'NMR', and, finally, (c) the AMBER99SB-STAR-ILDN force field,[22] which will be referred to as 'STAR'. The plain AMBER99SB force field has not been explicitly included in these calculations because results from the application of this force field to these peptides were already available from previous studies.[11,17,18]

2.2. System Preparation and Simulation Protocol. The production of the starting peptide structures in the fully extended state and solvation-ionization was performed with the program LEAP from the AMBER tools distribution.[23] For all simulations we used periodic boundary conditions and a cubic unit cell sufficiently large to guarantee a minimum separation between the PBC-related images of the peptides of at least 16 Å. We followed the dynamics of the peptides' folding simulations using the program NAMD[24] for a grand total of 21.7 μs using the TIP3P water model[25] and the force fields mentioned in the previous section (simulation lengths were 5.3, 5.0, and 5.0 μs for CLN025, 2.4, 2.0, and 2.0 μs for aLa). For the CLN025 peptide adaptive tempering[19] was applied as implemented in the program NAMD (Adaptive tempering is formally equivalent to a single-copy replica exchange folding simulation with a continuous temperature range. For the CLN025 simulations this temperature range was 300 to 500 K inclusive and was applied to the system through the Langevin thermostat, see below.).

The simulation protocol was the following. The systems were first energy minimized for 1000 conjugate gradient steps followed by a slow heating-up phase to a temperature of 320 K (with a temperature step of 20 K) over a period of 32 ps. Subsequently the systems were equilibrated for 10 ps under NpT conditions without any restraints, until the volume equilibrated. This was followed by the production NpT runs with the temperature and pressure controlled using the Nosè-Hoover Langevin dynamics and Langevin piston barostat control methods as implemented by the NAMD program, with adaptive tempering (for CLN025) applied through the Langevin thermostat, while the pressure was maintained at 1 atm. The Langevin damping coefficient was set to 1 ps^{-1}, and the piston's oscillation period was set to 200 fs, with a decay time of 100 fs. The production runs were performed with the impulse Verlet-I multiple time step integration algorithm as implemented by NAMD. The inner time step was 2 fs, short-range nonbonded interactions were calculated every one step, and long-range electrostatics interactions were calculated every

two timesteps using the particle mesh Ewald method with a grid spacing of approximately 1 Å and a tolerance of 10^{-6}. A cutoff for the van der Waals interactions was applied at 9 Å through a switching function, and SHAKE (with a tolerance of 10^{-8}) was used to restrain all bonds involving hydrogen atoms. Trajectories were obtained by saving the atomic coordinates of the whole systems every 0.8 ps.

2.3. Trajectory Analysis. The programs CARMA,[26] GRCARMA,[27] and Cluster5D[28] have been used for almost all of the analyses, including removal of overall rotations/translations, calculation of RMSDs from a chosen reference structure, calculation of the radius of gyration, calculation of the average structure (and of the atomic root mean squared fluctuations), production of PDB files from the trajectory, Cartesian space principal component analysis and corresponding cluster analysis, dihedral space principal component analysis and cluster analysis, calculation of the frame-to-frame RMSD matrices, calculation of similarity Q values, etc. Chemical shifts were calculated using the program SPARTA+[29] as previously described.[30] Secondary structure assignments were calculated with the program STRIDE.[31] All molecular graphics work and figure preparation were performed with the programs VMD,[32] RASTER3D,[33] and CARMA. Estimated NOE values were obtained from the trajectories using $\langle r^{-3} \rangle$-based averaging where r is the instantaneous distance between two selected protons from a trajectory. Reduced χ^2 values (for both NOEs and chemical shifts) were calculated as $\chi^2 = [\Sigma(S_{obs}-S_{calc})^2/\sigma^2]/\nu$ where S_{obs} and S_{calc} are the observed and calculated values of the quantity under examination, σ^2 is the corresponding estimated variance, and ν is the number of degrees of freedom. Please note that the estimated variances (σ^2) used for calculating the reduced χ^2 values were derived solely from the simulation-derived measurements and did *not* include estimates for additional sources of errors such as the experimental errors (generally very small compared with the simulation-derived uncertainties) or the errors introduced from the use of approximations such as those implemented in the SPARTA+ program. What this analysis implies, then, is that the χ^2 values quoted in this communication are safe, in the sense that they represent overestimates of the true quantities.

2.4. Extent of Sampling and Statistical Significance. Demonstrating statistical significance for the results obtained from a comparison between folding simulations appears to be a nightmare. To start with, it is almost certain that the sampling of the unfolded state is nowhere near convergence, and, thus, the comparison criteria must be based on derived quantities that avoid the unfolded state altogether. In addition, and as exemplified by the αLa peptide with its disordered C-terminus, the analysis must also be able to differentially focus on the potentially foldable substructures. Finally, even if the unfolded state and disordered substructures are excluded from the analyses, it is still highly improbable that all potentially accessible transiently stable peptide structures (corresponding to local minima of the peptide's energy landscape) will have been visited within the time scale of our simulations.

Having noted these problems, we believe that we can possibly establish statistical significance for our results by combining two different approaches. The first is to limit our comparisons on derived quantities − such as per residue secondary structure preferences and Q-T diagrams (discussed in the next section) − that can avoid the complications arising from the unfolded state and disordered substructures. The second, and more important, is based on an attempt to quantify

the extent of sampling of the simulations through a recently described probabilistic method which is based on the application of Good-Turing statistics to biomolecular trajectories.[34] The results obtained from this method are presented and discussed in the following paragraphs.

The application of Good-Turing statistics to biomolecular trajectory data produces an estimate of the probability of unobserved species (i.e. thus far unobserved protein/peptide structures) as a function of the RMSD of those unobserved structures from the structures that have already been observed in the simulation. For an example to clarify this, focus on the upper (black) curve from the first graph in Figure 1. The high
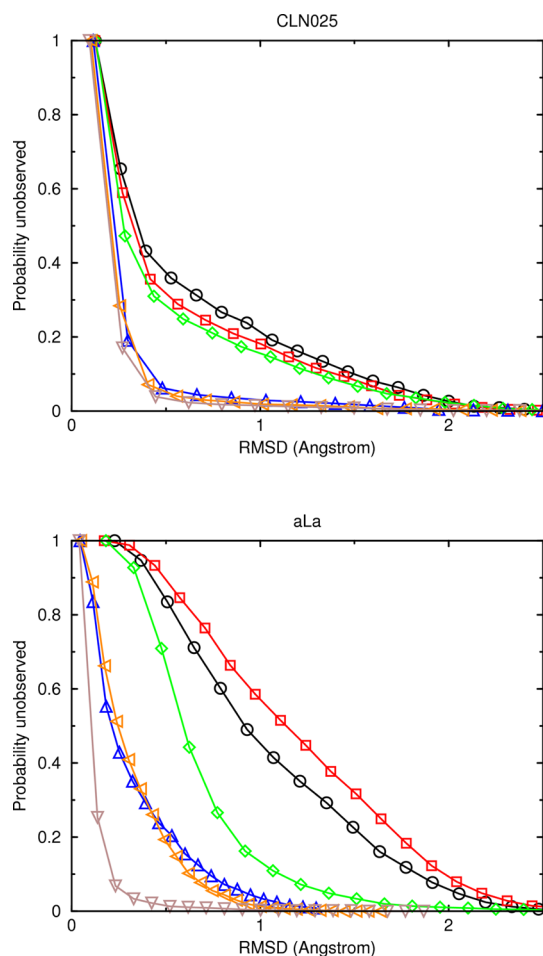
Figure 1. Extent of sampling and statistical significance. Results from the application of Good-Turing statistics to the folding trajectories of the CLN025 and αLa peptides. See section 2.4 for a discussion and analysis of these diagrams. For both diagrams the color coding is the following. For the upper three curves (of each diagram): black → STAR force field, red → ILDN, green → NMR. For the lower three curves: blue → STAR force field, orange → ILDN, brown → NMR.

probability of unobserved species for low RMSD values signifies the fact that if we were to continue the simulation it would be very probable to observe structures that although similar to some of the structures already observed, they would still differ slightly from them. As the RMSD (from the already observed structures) increases, the probability monotonically decreases. For a numerical example (and again referring to the black curve of the CLN025 diagram in Figure 1), we would expect that on average one out of five new (previously unobserved) structures

(corresponding to $P_{\text{unobserved}} = 0.20$) would differ by an RMSD of at least ~1.0 Å from the structures already observed or that one out of ten new structures ($P_{\text{unobserved}} = 0.10$) would differ by an RMSD of at least 1.5 Å, etc. It is the exact form of these graphs (and how quickly they reach low probability values) that signifies how well the trajectory has sampled the biomolecule's configurational space.

The upper diagram of Figure 1 shows the results from two sets of Good-Turing calculations as applied to the CLN025 peptide. The first set (corresponding to the upper three curves) is the direct application of the method to the three force fields studied in this paper and using for the analysis the whole of the trajectories (including the unfolded structures of the peptide). The lower three curves show the results from the same calculation but only using those peptide structures that are associated with low temperatures from adaptive tempering ($T < 340$ K). These low temperature structures correspond to relatively stable (from the simulation's point of view) peptide conformers. Taken together, these results show that although the whole of the peptide's configurational space (including unfolded structures) has not been adequately sampled (the upper three curves have relatively high $P_{\text{unobserved}}$ values even for large RMSD values), the stable peptide conformers demonstrate a much better sampling for all three force fields with negligible $P_{\text{unobserved}}$ values for RMSDs higher than ~1.5 Å. The implication of this analysis is that any differences observed between the stable peptide conformers (as obtained from the three force fields) are probably statistically significant.

The second diagram of Figure 1 shows the results from two sets of Good-Turing calculations as applied to the αLa peptide. The first set (corresponding to the upper three curves) is the direct application of the method to the three force fields studied in this paper and using for the analysis all peptide residues (including the disordered C-terminus). The lower three curves show the results from the same calculation but only using residues 2−7 of the peptide. As can be seen, the NMR force field (green and brown curves) differs appreciably in its behavior from the STAR and ILDN force fields. The reason for this difference will be discussed in section 3.2, but in summary the NMR force field overstabilizes an α-helical structure for these residues which leads to a significant reduction of the conformational space available to the αLa peptide. Taken together, these graphs show that although the whole of the peptide's configurational space (including its disordered residues) has not been adequately sampled (the upper three curves have relatively high $P_{\text{unobserved}}$ values even for large RMSD values), the residues that are known from the experiment to form the secondary structure (residues 2−7) demonstrate a much better sampling for all three force fields with negligible $P_{\text{unobserved}}$ values for RMSDs higher than ~1.2 Å. The implication of this analysis is that any differences observed in the secondary structure of residues 2−7 (as obtained from the three force fields) are probably statistically significant.

To conclude this section, it appears that the length of our simulations is probably sufficient for guaranteeing a reasonable sampling of the peptides' stable conformers, but it is not sufficient for studying differences over the whole of the peptides' folding landscapes. The implication of this analysis is that *not* all differences observed between the force fields can be considered significant. This will be discussed on a per analysis basis in the next section.
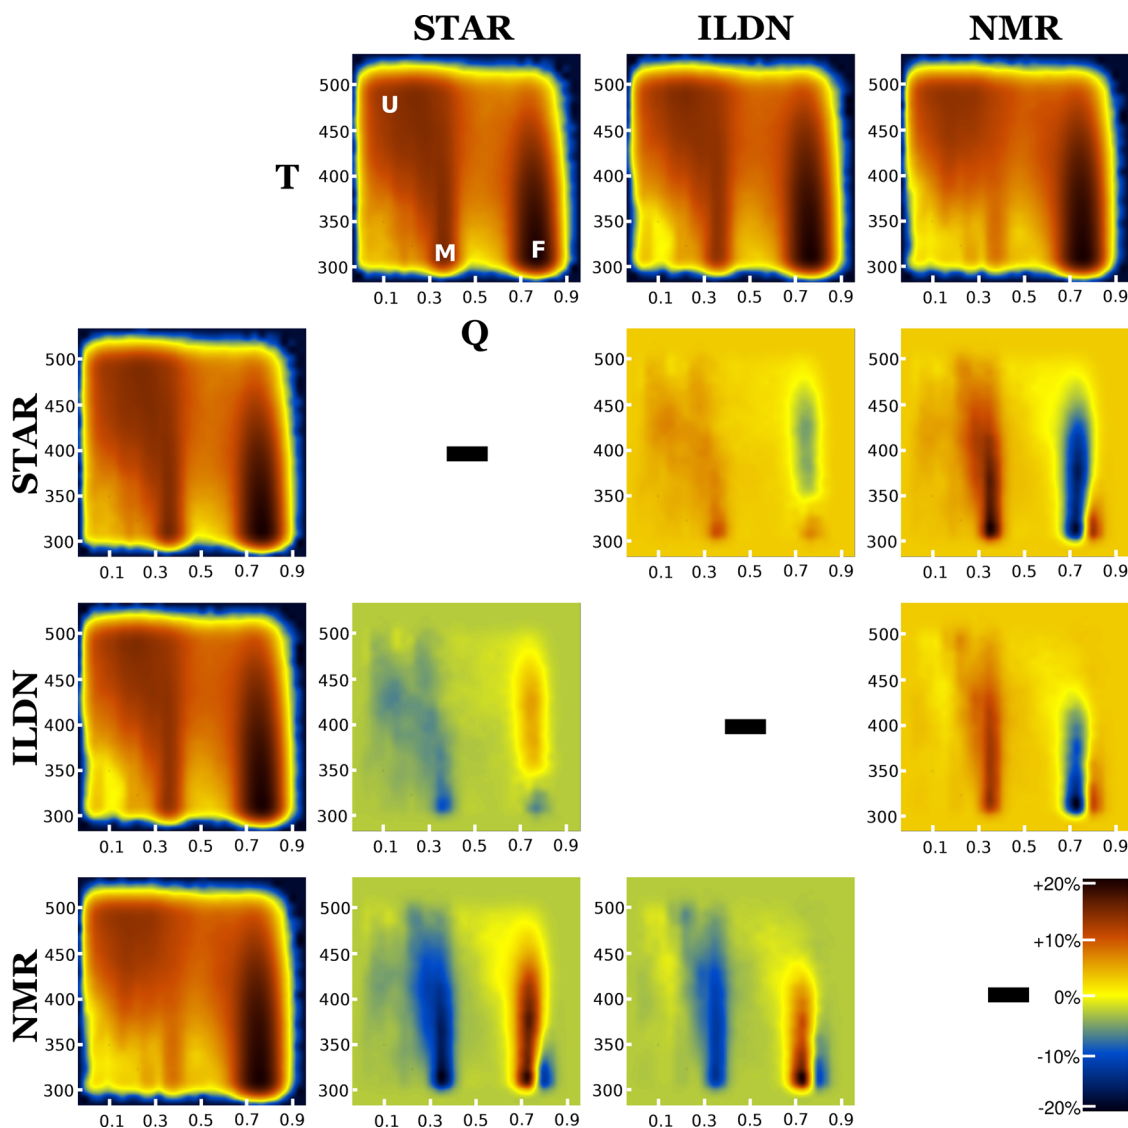
Figure 2. β-Hairpin peptide, difference Q-T diagrams. The top row and first column of this matrix are the Q-T diagrams of the corresponding force fields (see section 3.1 for a description). In the first diagram of the top row the unfolded, misfolded, and correctly folded peptide configurations are marked as U, M, and F, respectively. The contents of the matrix (the inner 3 by 3 graphs) are the differences between the respective force fields. The difference graphs were calculated as [(column force field) minus (row force field)]. The color bar at the lower right-hand corner expresses the observed differences in units of percent change compared to the maximum of the distributions (corresponding to the correctly folded peptide conformer marked as 'F').

3. RESULTS

3.1. For the Stably Folded β-Hairpin Peptide the Differences Are Relatively Minor.

In Figure 2 we compare the folding behavior of the CLN025 peptide for the three force fields using what will be referred to as Q-T diagrams.[35] In these diagrams (see the first graph in the top row of Figure 2) the vertical axis (T) corresponds to the adaptive tempering temperature from the simulation, the horizontal axis (Q) is the similarity index[36] to the experimentally known CLN025 structure (Q values range from 0.0 and 1.0, with a value of 1.0 corresponding to a structure identical with the experimentally determined), and hot colors in the graphs indicate high log density (i.e., that a large number of structures from the trajectory have the corresponding Q-T values). The calculations of Q values used in this communication are identical with the (q) values described in detail in the Supporting Information of ref 36. The high density peak near the lower

right-hand-side corner of these diagrams (marked as 'F' in Figure 2) corresponds to correctly folded peptide structures. The extent and density of the tail of this peak toward the high temperature regime is directly related to the thermal stability of the folded conformer. The peak marked as 'M' (at lower Q values) corresponds to the misfolded (1-offset β-hairpin) CLN025 structure. Finally, the extended area at the low Q/ high temperature part of the graph (marked as 'U') corresponds to unstable/unfolded peptide structures. Comparison between the Q-T diagrams for the three force fields (most easily seen in the first column of Figure 2) shows that in broad terms all three force fields perform well, and they all correctly fold the CLN025 peptide as expected. There are, however, significant differences between them as brought out by the inner 3 × 3 graphs of Figure 2 which depict the difference Q-T diagrams between the respective force fields. Examination of these difference maps shows that whereas the ILDN and STAR force fields are quite similar, the NMR variant appears to deviate

appreciably. In more detail, the main difference between the ILDN and STAR force fields is that ILDN slightly overstabilizes the native β-hairpin structure which (a) leads to a higher thermal stability of the folded structure [as indicated by the positive (red)/negative (blue) vertically aligned peaks at high Q values] and (b) reduces the population for the misfolded 1-offset hairpin (indicated by the difference peak at lower Q values). The NMR force field on the other hand appears to deviate significantly from both the ILDN and STAR variants. The most pronounced difference is that the Q values for the NMR-derived folded peptide structures are systematically lower than the Q values obtained from ILDN and STAR. This difference leads to the appearance of the characteristic side-by-side pattern of positive/negative peaks at high Q values. Surprisingly, and although the Q values are lower, the NMR force field appears to overstabilize the β-hairpin structure even more than the ILDN variant as can be deduced from the asymmetric length of the tails of the difference peaks toward the high temperature regime. The additional stabilization of the native β-hairpin structure demonstrated by the NMR variant further reduces the populations of the misfolded structures and gives rise to the pronounced difference peaks at medium Q values. These differences in the stabilization of the native β-hairpin structure by the different force fields can more readily be seen in Supporting Information Figures S1 and S2 which compare the melting curves and the fraction of folded structures vs temperature for the three force fields.

The discussion in the previous paragraph established that there are indeed statistically significant differences between the three force fields. The systematic difference observed in the preferred Q values between the ILDN/STAR and NMR variants suggests that the ILDN and STAR force fields reproduce more closely the experimentally determined CLN025 structure. To confirm that this is indeed the case we have compared the representative folded peptide structures obtained from the three force fields with the experimental structure. The RMSDs (after least-squares superposition) between the experimental and the ILDN, STAR, and NMR representative structures were found to be 0.75, 0.77, and 0.81 Å respectively for the Cα atoms and 0.87, 0.92, and 0.97 Å for all backbone atoms. These structural differences do establish that STAR and ILDN reproduce more closely the experimental structure but do not differentiate which one of the three force field variants is closer to the raw experimental data. We have addressed this issue by calculating the reduced χ^2 value (goodness-of-fit) between the experimentally determined NOEs (as deposited with the PDB entry 2rvd) and the values expected from the distributions of interproton distances obtained from the three simulations. The reduced χ^2 values were found to be 2.4 for the STAR variant, 2.8 for ILDN, and 3.3 for the NMR force field, indicating the better agreement of STAR with the experiment [The quoted χ^2 values were calculated from all structures whose associated adaptive tempering temperature was less than 320 K. Essentially identical results were obtained with all other temperature ranges that we have tested]. Since for the CLN025 peptide chemical shifts have also been deposited with the Protein Data Bank (PDB entry 2rvd), we have also calculated the reduced χ^2 values between the experimental and simulation-derived chemical shifts (chemical shifts were calculated from the simulations using the program SPARTA+[29] as previously described[30]). For the HN chemical shifts the reduced χ^2 values were found to be 1.129 for the STAR variant, 1.134 for ILDN,

and 1.158 for the NMR force field, in agreement with the results obtained from the NOEs. Similar results were also obtained by examining the HA chemical shifts which gave reduced χ^2 values of 0.788 for the STAR variant, 0.796 for ILDN, and 0.798 for the NMR force field. The significantly lower χ^2 values obtained from the chemical shifts (compared with the comparison of NOEs) are mostly the result of the relatively high standard deviations of the simulation-derived values as can be seen in Supporting Information Figure S3. We close this paragraph by noting that the lower melting temperature predicted by the STAR variant (see Supporting Information Figures S1 and S2) is also in better agreement with the experimental thermal stability data available for the CLN025 peptide.

3.2. For the Less Stable Helical Peptide the NMR Variant Deviates Significantly. Figure 3 shows the STRIDE-
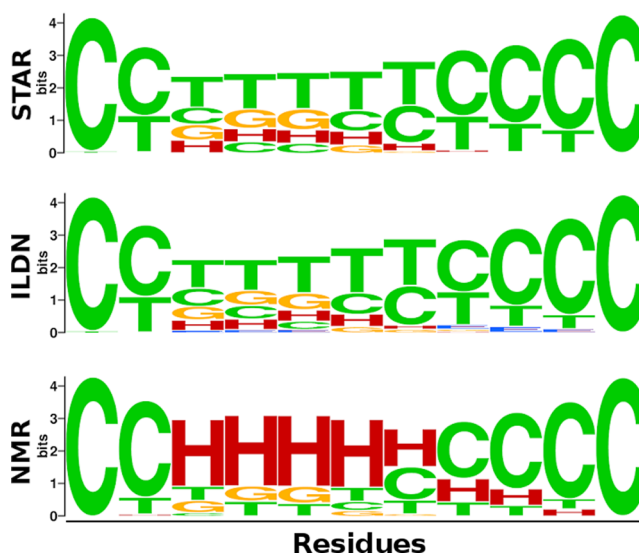
Figure 3. Helical peptide, secondary structure preferences. This diagram shows WebLogo-like representations of the STRIDE-assigned per residue secondary structure preferences for the αLa peptide and for each of the three force fields examined. The symbols are C → coil (green), T → turn (green), G → 3$_{10}$ helix (orange), H → α-helix (red), E → extended beta structure (blue).

derived[31] per residue secondary structure preferences for the αLa peptide and for each of the three force fields studied. Again, and as observed for the β-hairpin peptide, the STAR and ILDN variants appear to be quite similar to relatively minor differences, such as the sightly higher 3$_{10}$- and α-helical populations for STAR (most easily seen for residue 4), and the presence of a very minor β-population for the ILDN force field. Both STAR and ILDN predict a marginally stable helical population for residues 3−6, amounting to approximately 40% of the trajectories' total length, and an almost completely disordered C-terminus (residues 7−11). In contrast, the NMR variant overstabilizes an almost purely α-helical population at the expense of both the 3$_{10}$-helical and disordered conformations. The overstabilization of the α-helical structure in the case of NMR is such that the helical population reaches approximately 80% of the trajectory's total length for residues 4−5 and extends with significant helicity out to residues 7−10. These pronounced differences between the secondary structure preferences can also be seen in Supporting Information Figure S4 which shows a direct comparison between the Ramachan-

dran plots of the STAR and NMR variants along with a difference Ramachandran plot.

The preceding analysis established the presence of appreciable differences between the force fields but has not addressed the issue of how the three variants perform in comparison with the experimental data. The αLa peptide is known from several independent CD and NMR experiments[15,16] to form a marginally stable helical 3_{10}-like structure with an estimated population for its foldable part of approximately 40% as deduced from (a) the $[\Theta]_{215}/[\Theta]_{198}$ and $[\Theta]_{208}/[\Theta]_{198}$ ellipticity ratios,[16] (b) by analyzing the integrated cross peak intensity of NOESY spectra,[16] (c) by calculating the peptide's helical fraction from the molar ellipticity at 222 nm,[15] (d) by comparing the CD-derived helical fraction between the peptide in aqueous solution and the peptide in 30% TFE,[15] and, finally, (e) by comparing — on a per residue basis — the number and magnitude of deviations of the $C^{\alpha}H$ chemical shifts from random coil values for the peptide in aqueous solution and the peptide in 30% TFE.[15] Additionally, the STRIDE-derived secondary structure assignments of the NMR-derived peptide structure conformers (PDB entry 1CB3) show a mixture of 3_{10}-helix and α-helix for residues 3−6, with the majority of the assignments being 3_{10}-helical. These results clearly indicate that the STAR and ILDN variants are in much better agreement with the experiment than the NMR force field both in predicting the peptide's marginal stability and in preferentially populating a 3_{10}-helical conformation (instead of a pure α-helical conformation as is the case with the NMR force field). To put these observations in numbers we have also calculated for each of the three force fields the number of upper bound violations of the deposited NOE restraints, as well as the values of the corresponding average violations (NOE data were obtained from the PDB entry 1cb3). For a grand total of 138 restraints, the STAR force field gave 17 upper bound violations with an average violation of 0.10 Å, the ILDN variant gave 18 violations with an average of 0.12 Å, and NMR gave 20 violations with an average of 0.15 Å, in good agreement with the results obtained from the secondary structure analysis. Please note that a comparison with experimentally determined chemical shifts (as previously described for the CLN025 peptide) can not be performed in the case of αLa because no such data have been deposited for this peptide.

In summary, for the less stable helical peptide the NMR force field clearly deviates from the STAR and ILDN variants by overstabilizing and α-helical structure. The STAR and ILDN variants both reproduce closely the experimentally available data for the peptide, with the STAR force field performing detectably better than the ILDN variant.

4. DISCUSSION

We believe that we have convincingly established the sensitivity of folding simulations to even relatively limited variation within the same force field subfamily. Using two completely different peptide systems, we have demonstrated that three different variants of the AMBER99SB force field show statistically significant differences even in the case of a fast-folding and very stable β-hairpin peptide. For the marginally stable 3_{10}-helical peptide the differences were so pronounced that even a simple enumeration of secondary structure preferences sufficed for differentiating between the three variants. Clearly, and as a cursory examination of Figure 3 and of Supporting Information Figure S1 shows, the word 'variant' may be a poor choice

considering the amount of differences between the force fields examined.

In a sense, our results reinforce what is common knowledge in the field: development of a balanced biomolecular force field is an extremely difficult and delicate procedure. In this respect, the consistently better performance of the STAR force field serves as a testimony to the care taken during its development to avoid overfitting[22] and suggests that maybe we have reached the stage where we should substitute the expression 'force field development' with the expression 'force field refinement'.

■ ASSOCIATED CONTENT

ⓢ Supporting Information

The Supporting Information is available free of charge on the ACS Publications website at DOI: 10.1021/acs.jcim.6b00493.

Figures S1−S4 and Table S1 (PDF)

■ AUTHOR INFORMATION

Corresponding Author

*Phone: +30-25510-30620. Fax: +30-25510-30620. E-mail: glykos@mbg.duth.gr. Web address: http://utopia.duth.gr/glykos/.

Notes

The authors declare no competing financial interest.

■ REFERENCES

(1) Shaw, D. E.; Maragakis, P.; Lindorff-Larsen, K.; Piana, S.; Dror, R. O.; Eastwood, M. P.; Bank, J. A.; Jumper, J. M.; Salmon, J. K.; Shan, Y.; Wriggers, W. Atomic-Level Characterization of the Structural Dynamics of Proteins. *Science* **2010**, *330*, 341−346.

(2) Gnanakaran, S.; García, A. E. Validation of an All-Atom Protein Force Field: From Dipeptides to Larger Peptides. *J. Phys. Chem. B* **2003**, *107*, 12555−12557.

(3) Daggett, V.; Kollman, P. A.; Kuntz, I. D. A Molecular Dynamics Simulation of Polyalanine: An Analysis of Equilibrium Motions and Helix-Coil Transitions. *Biopolymers* **1991**, *31*, 1115−1134.

(4) Gnanakaran, S.; García, A. E. Helix-Coil Transition of Alanine Peptides in Water: Force Field Dependence on the Folded and Unfolded Structures. *Proteins: Struct., Funct., Genet.* **2005**, *59*, 773−782.

(5) Matthes, D.; de Groot, B. L. Secondary Structure Propensities in Peptide Folding Simulations: A Systematic Comparison of Molecular Mechanics Interaction Schemes. *Biophys. J.* **2009**, *97*, 599−608.

(6) Lindorff-Larsen, K.; Maragakis, P.; Piana, S.; Eastwood, M. P.; Dror, R. O.; Shaw, D. E. Systematic Validation of Protein Force Fields against Experimental Data. *PLoS One* **2012**, *7*, e32131.

(7) Piana, S.; Klepeis, J. L.; Shaw, D. E. Assessing the Accuracy of Physical Models Used in Protein-Folding Simulations: Quantitative Evidence from Long Molecular Dynamics Simulations. *Curr. Opin. Struct. Biol.* **2014**, *24*, 98−105.

(8) Kaur, H.; Sasidhar, Y. U. For the Sequence YKGQ, the Turn and Extended Conformational Forms Are Separated by Small Barriers and the Turn Propensity Persists Even at High Temperatures: Implications for Protein Folding. *J. Phys. Chem. B* **2012**, *116*, 3850−3860.

(9) Lindorff-Larsen, K.; Piana, S.; Dror, R. O.; Shaw, D. E. How fast-folding proteins fold. *Science* **2011**, *334*, 517−520.

(10) Best, R. B. Atomistic Molecular Simulations of Protein Folding. *Curr. Opin. Struct. Biol.* **2012**, *22*, 52−61.

(11) Patapati, K. K.; Glykos, N. M. Three Force Fields' Views of the 3_{10} Helix. *Biophys. J.* **2011**, *101*, 1766−1771.

(12) Georgoulia, P. S.; Glykos, N. M. Using J-Coupling Constants for Force Field Validation: Application to Hepta-alanine. *J. Phys. Chem. B* **2011**, *115*, 15221−15227.

(13) Hornak, V.; Abel, R.; Okur, A.; Strockbine, B.; Roitberg, A.; Simmerling, C. Comparison of Multiple Amber Force Fields and

Development of Improved Protein Backbone Parameters. *Proteins: Struct., Funct., Genet.* 2006, 65, 712−725.

(14) Wickstrom, L.; Okur, A.; Simmerling, C. Evaluating the Performance of the ff99SB Force Field Based on NMR Scalar Coupling Data. *Biophys. J.* 2009, 97, 853−856.

(15) Demarest, S. J.; Hua, Y.; Raleigh, D. P. Local Interactions Drive the Formation of Non-Native Structure in the Denatured State of Human α-Lactalbumin: a High Resolution Structural Characterization of a Peptide Model in Aqueous Solution. *Biochemistry* 1999, 38, 7380−7387.

(16) Araki, M.; Tamura, A. Transformation of an α-Helix Peptide into a β-Hairpin Induced by Addition of a Fragment Results in Creation of a Coexisting State. *Proteins: Struct., Funct., Genet.* 2007, 66, 860−868.

(17) Patapati, K. K.; Glykos, N. M. Order through Disorder: Hyper-Mobile C-Terminal Residues Stabilize the Folded State of a Helical Peptide. A Molecular Dynamics Study. *PLoS One* 2010, 5, e15290.

(18) Honda, S.; Akiba, T.; Kato, Y. S.; Sawada, Y.; Sekijima, M.; Ishimura, M.; Ooishi, A.; Watanabe, H.; Odahara, T.; Harata, K. Crystal Structure of a ten-amino acid protein. *J. Am. Chem. Soc.* 2008, 130, 15327−31.

(19) Zhang, C.; Ma, J. Enhanced Sampling and Applications in Protein Folding in Explicit Solvent. *J. Chem. Phys.* 2010, 132, 244101.

(20) Lindorff-Larsen, K.; Piana, S.; Palmo, K.; Maragakis, P.; Klepeis, J. L.; Dror, R. O.; Shaw, D. E. Improved Side-Chain Torsion Potentials for the Amber ff99SB Protein Force Field. *Proteins: Struct., Funct., Genet.* 2010, 78, 1950−1958.

(21) Li, D. W.; Bruschweiler, R. NMR-based Protein Potentials. *Angew. Chem., Int. Ed.* 2010, 49, 6778−6780.

(22) Best, R. B.; Hummer, G. Optimized Molecular Dynamics Force Fields Applied to the Helix-Coil Transition of Polypeptides. *J. Phys. Chem. B* 2009, 113, 9004−9015.

(23) Case, D. A.; Cheatham, T. E., III; Darden, T.; Gohlke, H.; Luo, R.; Merz, K. M., Jr.; Onufriev, A.; Simmerling, C.; Wang, B.; Woods, R. J. The Amber Biomolecular Simulation Programs. *J. Comput. Chem.* 2005, 26, 1668−1688.

(24) Kale, L.; Skeel, R.; Bhandarkar, M.; Brunner, R.; Gursoy, A.; Krawetz, N.; Phillips, J.; Shinozaki, A.; Varadarajan, K.; Schulten, K. NAMD2: Greater Scalability for Parallel Molecular Dynamics. *J. Comput. Phys.* 1999, 151, 283−312.

(25) Jorgensen, W. L.; Chandrasekhar, J.; Madura, J. D.; Impey, R. W.; Klein, M. L. Comparison of Simple Potential Functions for Simulating Liquid Water. *J. Chem. Phys.* 1983, 79, 926−935.

(26) Glykos, N. M. CARMA: A Molecular Dynamics Analysis Program. *J. Comput. Chem.* 2006, 27, 1765−1768.

(27) Koukos, P. I.; Glykos, N. M. grcarma: A Fully Automated Task-Oriented Interface for the Analysis of Molecular Dynamics Trajectories. *J. Comput. Chem.* 2013, 34, 2310−2312.

(28) Baltzis, A. S.; Koukos, P. I.; Glykos, N. M. Clustering of Molecular Dynamics Trajectories via Peak-Picking in Multidimensional PCA-derived Distributions. *arXiv.* 2015, 1512.04024 [q-bio.BM].

(29) Shen, Y.; Bax, A. SPARTA+: a Modest Improvement in Empirical NMR Chemical Shift Prediction by Means of an Artificial Neural Network. *J. Biomol. NMR* 2010, 48, 13−22.

(30) Baltzis, A. S.; Glykos, N. M. Characterizing a Partially Ordered Miniprotein through Folding Molecular Dynamics Simulations: Comparison with the Experimental Data. *Protein Sci.* 2016, 25, 587−596.

(31) Frishman, D.; Argos, P. Knowledge-Based Protein Secondary Structure Assignment. *Proteins: Struct., Funct., Genet.* 1995, 23, 566−579.

(32) Humphrey, W.; Dalke, A.; Schulten, K. VMD—Visual Molecular Dynamics. *J. Mol. Graphics* 1996, 14, 33−38.

(33) Merritt, E. A.; Bacon, D. J. Raster3D Photorealistic Molecular Graphics. *Methods Enzymol.* 1997, 277, 505−524.

(34) Koukos, P. I.; Glykos, N. M. On the Application of Good-Turing Statistics to Quantify Convergence of Biomolecular Simulations. *J. Chem. Inf. Model.* 2014, 54, 209−217.

(35) Koukos, P. I.; Glykos, N. M. Folding Molecular Dynamics Simulations Accurately Predict the Effect of Mutations on the Stability and Structure of a Vammin-Derived Peptide. *J. Phys. Chem. B* 2014, 118, 10076−10084.

(36) Cho, S. S.; Levy, Y.; Wolynes, P. G. P versus Q: Structural Reaction Coordinates Capture Protein Folding on Smooth Landscapes. *Proc. Natl. Acad. Sci. U. S. A.* 2006, 103, 586−591.