Supplementary Figures

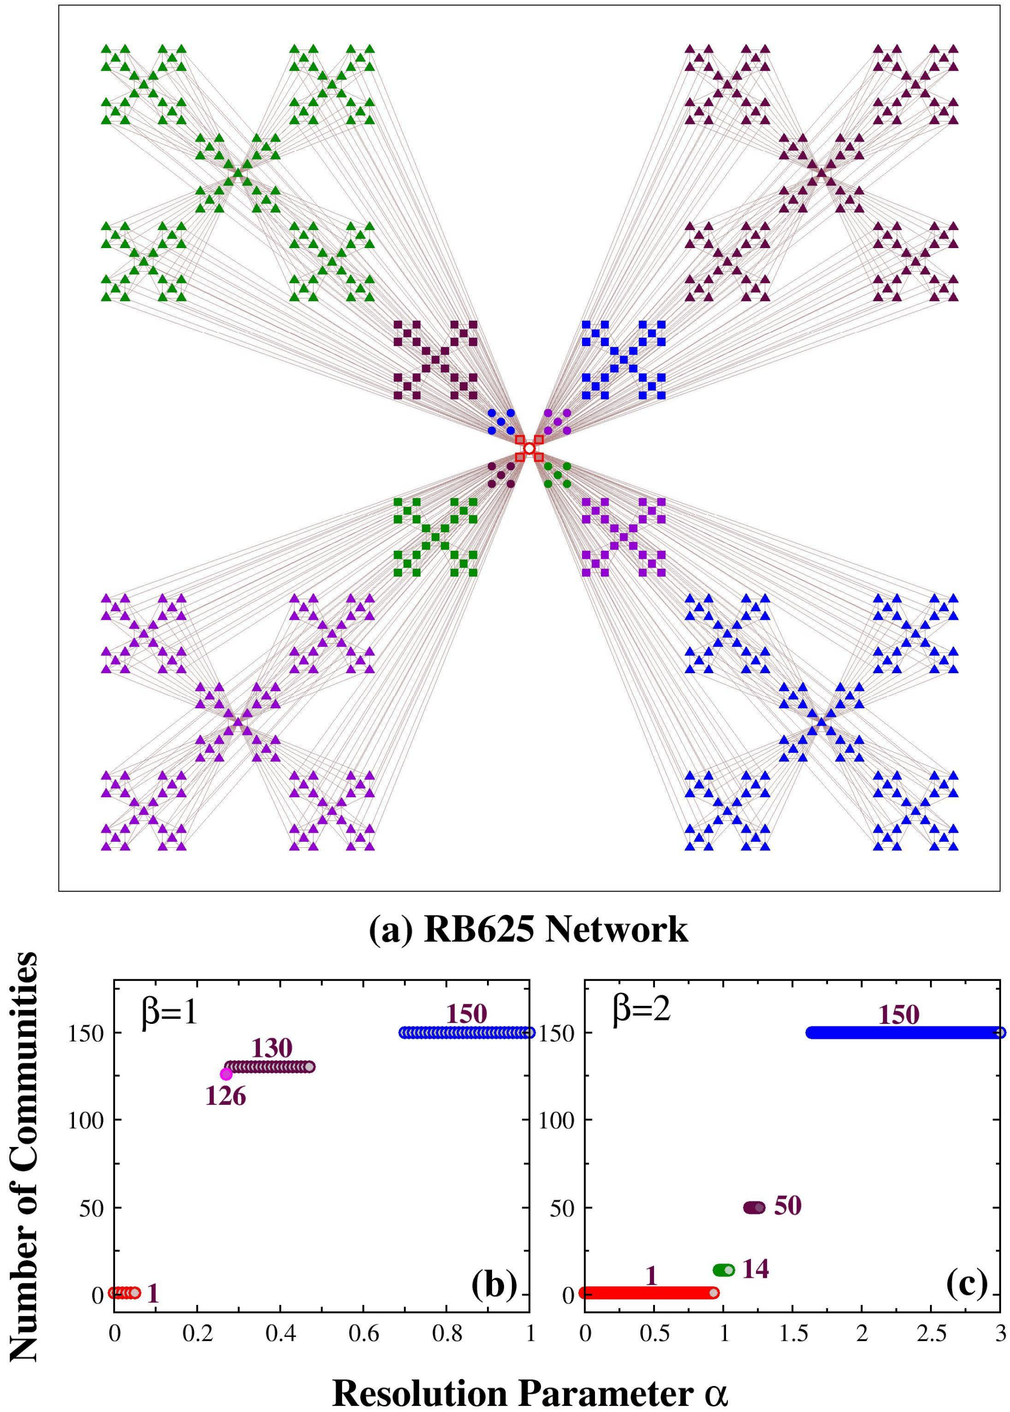

Supplementary Figure S1. Communities in the RB625 network. (a) Structure of the RB625 network. Here, by different shapes and colors of nodes, we show a division of the network into 14 communities on the 3rd community level. (b) and (c): Plateaus detected for the RB625 network by our method with $\beta = 1$ and 2 within 1000 realizations of calculation at each resolution.

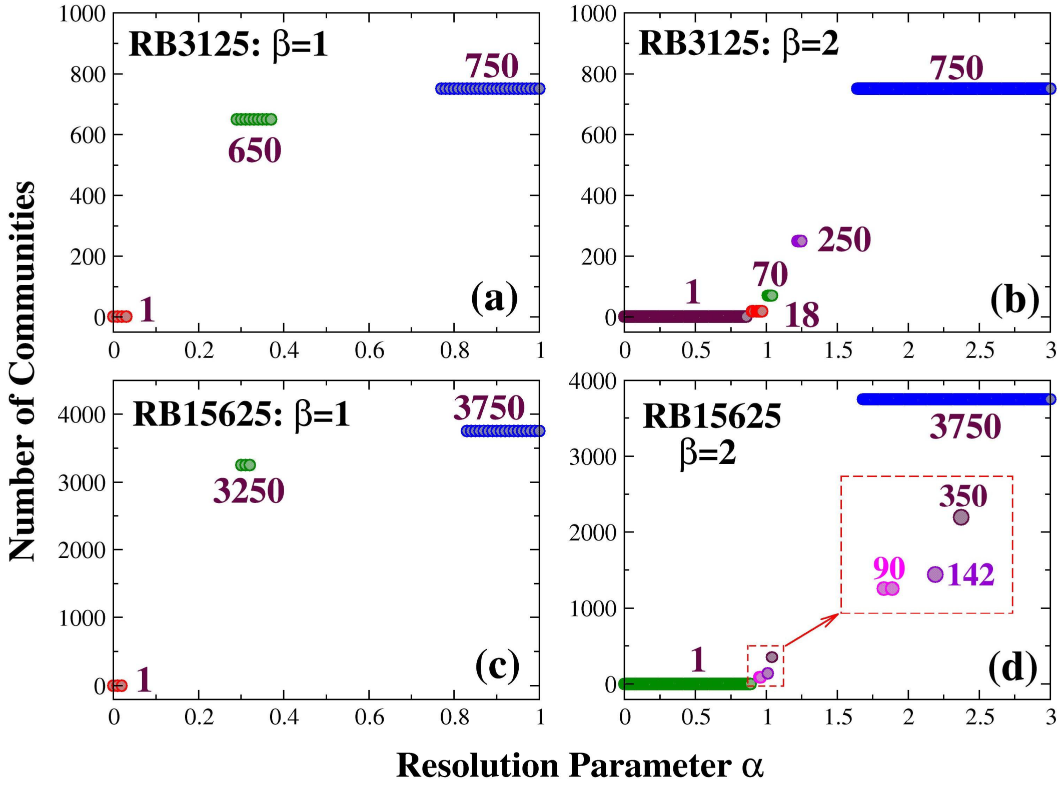

Supplementary Figure S2. Plateaus detected for the RB3125 and RB15625 networks with $\beta = 1$ and 2 within 1000 realizations of calculation at each resolution. Subfigures (a–c) are perfectly consistent with our discussion in the main text. Yet in (d), for the RB15625 network, our detection starts to deviate from perfection: (1) the 2nd and 5th community levels (containing 1250 and 22 communities, respectively) are undetectable, and (2) a small plateau containing 142 communities emerges, which turns out to be a hybrid of different levels of communities: each of the four peripheral RB3125 units is divided into 18 communities (as on the 4th community level of an RB3125 network), while the central RB3125 unit is divided into 70 communities (as on the 3rd community level thereof).

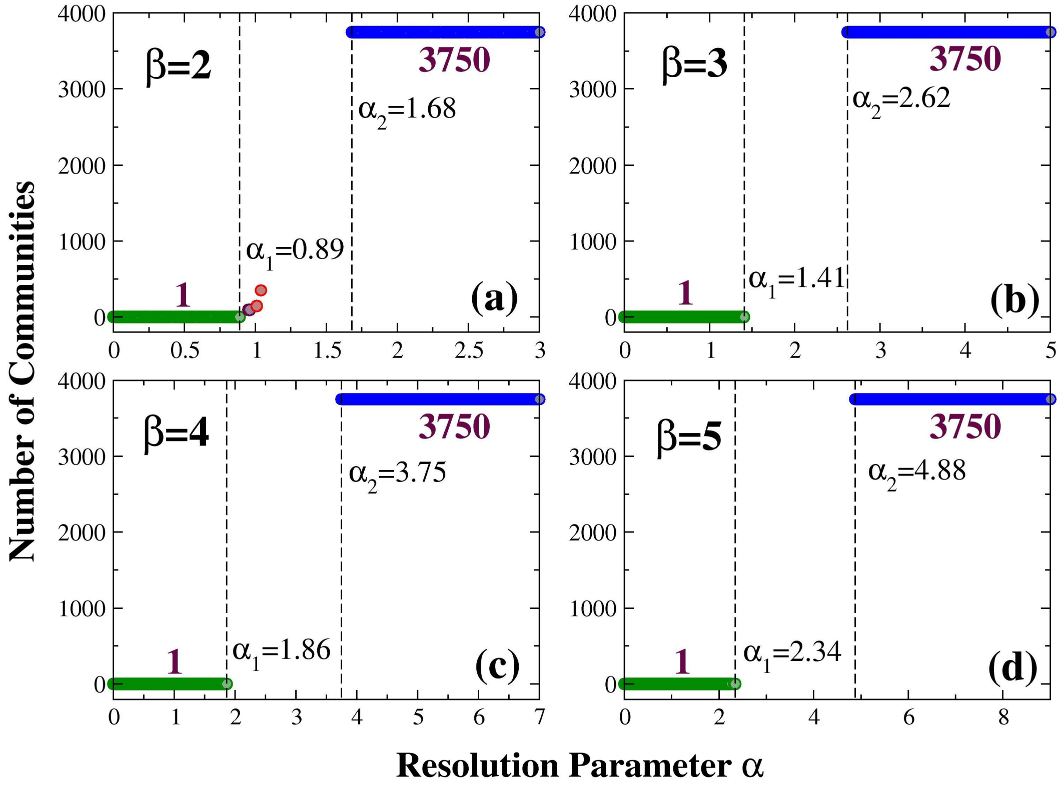

Supplementary Figure S3. Plateaus for the RB15625 network detected by (a) $\beta = 2$, (b) $\beta = 3$, (c) $\beta = 4$ and (d) $\beta = 5$ in 1000 realizations of calculation at each resolution. α_1 and α_2 in the figure indicate the upper and lower bounds of the resolution scales of the highest and lowest community levels; all intermediate community levels are limited within the range $\alpha_1 < \alpha < \alpha_2$ (plateaus not shown). It can be observed that this range does not increase linearly with β; increasing β does not help to detect more than 5 levels of communities for the RB networks.

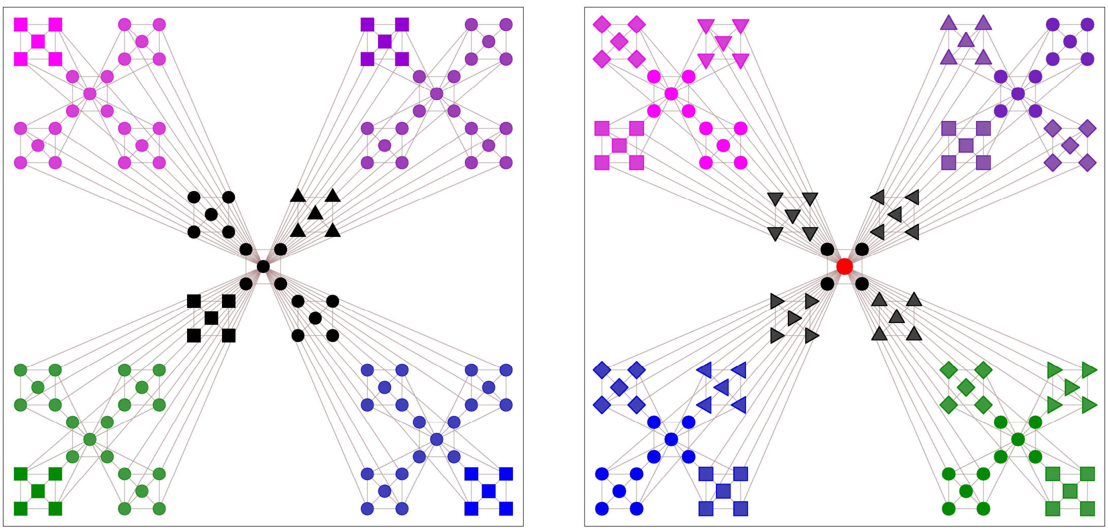

(a) Modularity Q of Newman **(b) Infomap**

Supplementary Figure S4. Communities in an RB125 network detected by (a) the standard modularity function Q proposed by Newman [48] and (b) Infomap [37]. (a) Modularity Q divides the network into 11 communities *in a random manner*: two RB5 units (black squares and triangles in (a)) are *randomly* chosen as two individual communities from the central RB25 unit, and the rest three

RB5 units are integrated as a third community (black circles in (a)). Similarly, each peripheral RB25 unit is divided into two communities by *randomly* choosing one of its five RB5 units as an individual community but integrating the rest four as another community. (b) Infomap divides the network into 22 communities: the central RB25 unit is divided into six communities as in Figure 2 (c) in the main text, while each of the peripheral RB25 units is divided into four communities by *randomly* combing one of its four peripheral RB5 units with the central unit, and having the rest three as three individual communities. Since both these divisions emerge with randomness, due to the symmetry of the RB networks, neither of these divisions is best-and-unique, nor are they stable solutions for community structures within the RB networks.

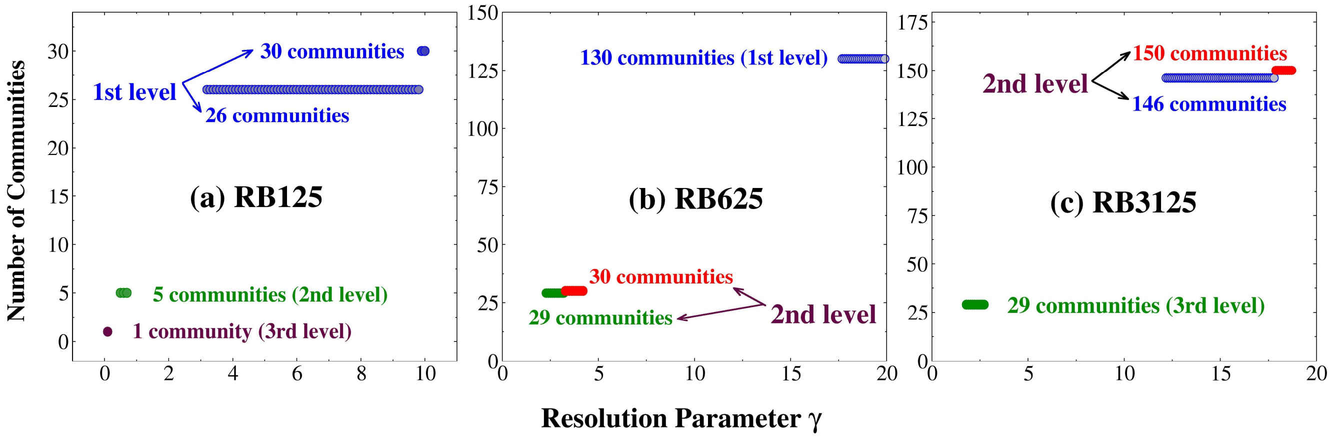

Supplementary Figure S5. Plateaus identified by a generalized version of modularity Q proposed by Reichardt and Bornholdt [33], i.e., $Q_\gamma = \sum_{g}\left[\frac{k_{in}^{g}}{2M} - \gamma\left(\frac{k^{g}}{2M}\right)^2\right]$. Here the plateaus are defined with the stringency as we suggested in section 2.3, which requires "one plateau one topology" and each data point representing a "best-and-unique" solution. With such a stringency, (a) for an RB125 network, Q_γ detects all three community levels: on the first (lowest) level, it detects both the "robust" division (30 communities) and a variant division (26 communities) which we have both discussed in section 3.1 of the main text, while on the second level, it detects only the "natural" division. (b) and (c): For larger RB networks such as RB625 and RB3125, Q_γ detects no more than two community levels for each of them; other levels are all undetectable. Note that in (c), we have extended our calculation until parameter γ is as large as 60, but still cannot detect a stable division of the first community level with Q_γ.

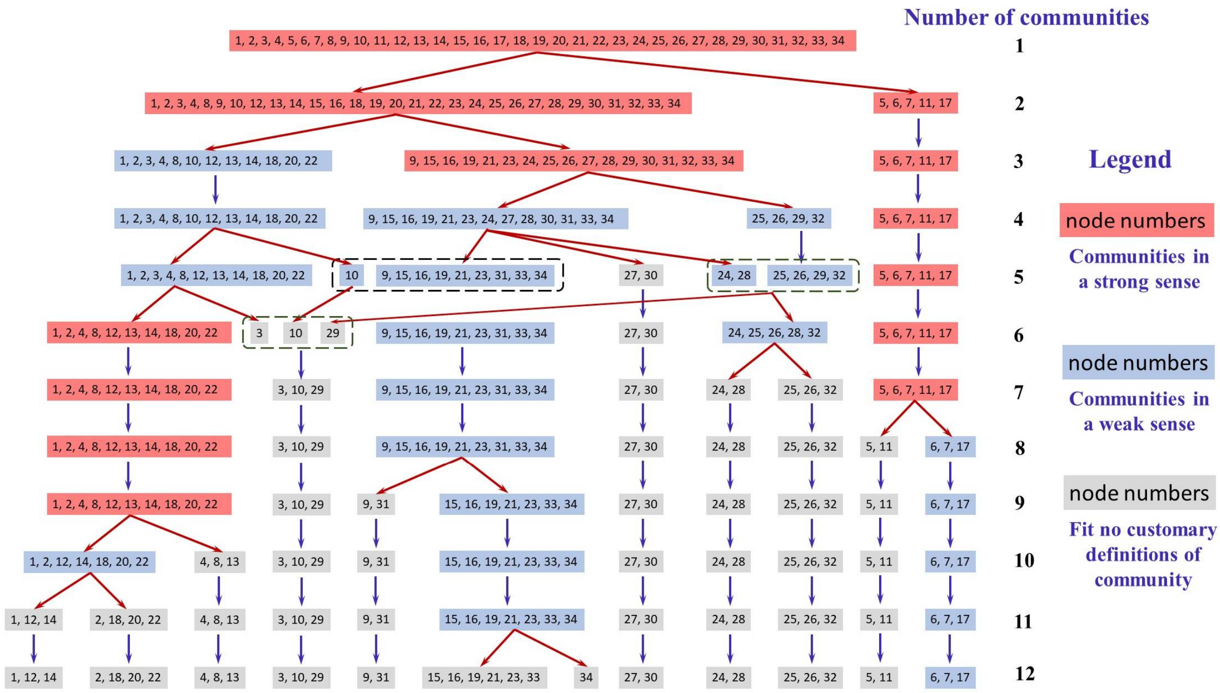

Supplementary Figure S6. Multi-level community structures within the karate club network detected by our method. Different levels of communities roughly exhibit a hierarchical structure, with minor reassembling of communities (indicated by the dashed rectangles) between a few of the neighboring levels. We distinguish communities defined with different stringencies by different fill colors: red indicates communities defined in a strong sense, blue in a weak sense, while grey fits no customary definitions of community.

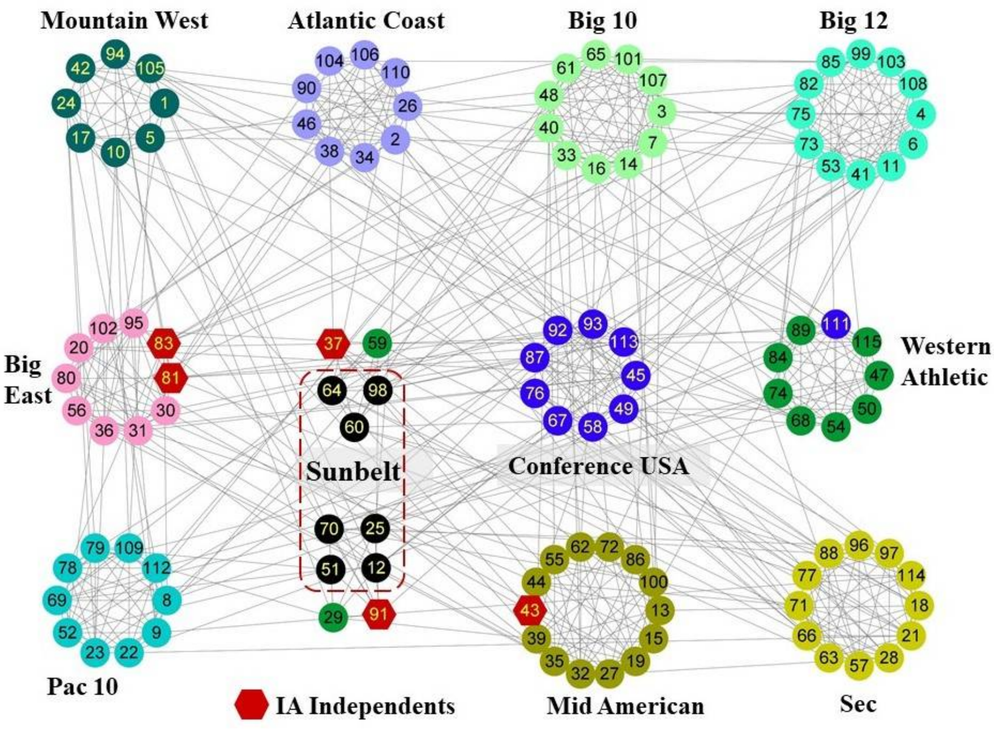

Supplementary Figure S7. Twelve communities detected in the American college football network by our method, in comparison with the metadata on conference assignments recorded in Figure 5 of [2]. There are some observable discrepancies between our detection and the metadata: (1) Conference Sunbelt (black nodes enclosed in the red dashed box) is split into two parts, whose members seldom played any games with the members of the other part. (2) Node 111 (Texas Christian) was recorded as a member of "Conference USA," but it did not play even one single game with any other teams of the same conference—instead it played quite some games with teams in conference Western Athletic and is then assigned to the community of the latter. (3) For the same reason, node 29 (Boise State) is assigned to one of the communities of Sunbelt instead of Western Athletic. Girvan's and Newman's division agrees with ours on all the above (1)–(3); the only difference still lies in the assignment for node 37 (Central Florida). In 2010, Evans pointed out there was a serious error in the above metadata [66]: the conference assignments were collected during the 2001 season, not the 2000 season! The proof is, conference Big West existed for football till 2000 while conference Sun Belt was only started in 2001. With the metadata corrected in [66], both our division and Girvan–Newman's perfectly recreate members of all conferences (see Figure 8 in the main text). As argued by the authors of [55], human errors can render the metadata irrelevant to the network structure.

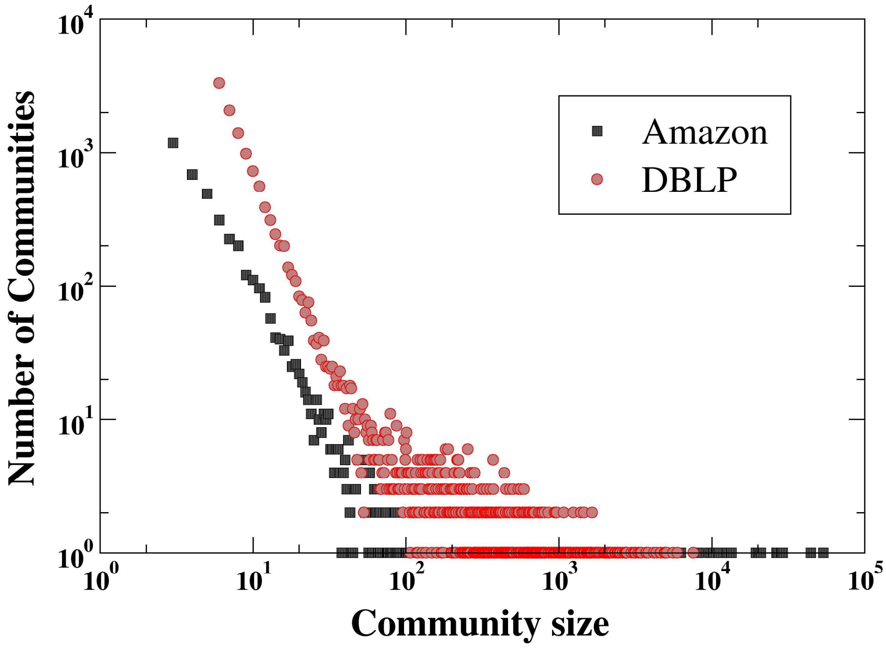

Supplementary Figure S8. Community size distributions for the Amazon and DBLP networks derived from the ground truth communities suggested for these two networks by [67]. These two networks have similar sizes (around 300,000 nodes), and both distributions are power-law of similar exponent. DBLP contains more communities, and Amazon has communities of larger sizes. When detecting communities in these networks with the multiresolution modularity Q_γ, it is obvious that DBLP would require a larger value of γ than Amazon for the ground truth communities being detected. Yet without knowing the community sizes beforehand, it is not possible to estimate even roughly the resolution scales within which these ground truth communities should be detected. Sometimes one has to carry out a huge amount of calculations with γ varying over several orders of magnitude to locate the optimized resolution scale within which the expected community structure can be recovered.