WHAT IS CLAIMED IS:

1. A recommendation system for selecting items to recommend, the recommendation system comprising:

a candidate generator operative to generate a ranked set of candidate items to recommend to a target user;

a probabilistic scorer operative to probabilistically vary rankings of at least some of the candidate items in the ranked set to generate a probabilistically-varied ranked set of candidate items; and

a candidate filter operative to select, from said probabilistically-varied ranked set, a subset of items to present to the user.

2. The recommendation system of Claim 1, wherein the degree to which the probabilistic scorer probabilistically varies the ranking of a candidate item depends on a number of identified reasons for recommending the candidate item to the target user.

3. The recommendation system of Claim 2, wherein probabilistic scorer causes candidate items having relatively low numbers of reasons to be probabilistically varied to a greater degree than candidate items having relatively high numbers of reasons.

4. A computer-implemented method of selecting items to recommend, the method comprising:

generating a ranked set of candidate items to recommend to a target user;

probabilistically varying rankings of at least some of the candidate items in the ranked set to generate a probabilistically-varied ranked set of candidate items; and

selecting, from said probabilistically-varied ranked set, a subset of items to present to the user.

5. The computer-implemented method of Claim 4, wherein selecting, from said probabilistically-varied ranked set, a subset of items to present to the user comprises selecting a most-highly ranked subset of items.

6. The computer-implemented method of Claim 4, wherein at least some of the candidate items are associated with one or more identified reasons for being recommended to the target user.

7. The computer-implemented method of Claim 6, wherein the degree to which the ranking of a candidate item is probabilistically varied depends on the number of identified reasons for recommending the candidate item to the target user.

8. The computer-implemented method of Claim 7, wherein probabilistically varying the rankings comprises causing candidate items having relatively low numbers of reasons to be probabilistically varied to a greater degree than candidate items having relatively high numbers of reasons.

9. The computer-implemented method of Claim 4, wherein probabilistically varying the rankings comprises causing candidate items having relatively low category scores to be probabilistically varied to a greater degree than candidate items having relatively high category scores.

10. The computer-implemented method of Claim 4, wherein probabilistically varying the rankings comprises generating values for at least a portion of the candidate items and adding probabilistic noise to the values.

11. The computer-implemented method of Claim 10, wherein adding probabilistic noise to the values comprises generating a probability distribution corresponding to each value, wherein the characteristics of the probability distribution depend on the corresponding value.

12. The computer-implemented method of Claim 11, further comprising sampling values from each probability distribution to generate a score for said at least a portion of the candidate items, wherein the score is responsive to the sampled value.

13. The computer-implemented method of Claim 4, wherein probabilistically varying rankings of at least some of the candidate items in the ranked set further comprises probabilistically varying the rankings to achieve an overlap of 5 to 20 percent of recommended items between recommendation generation events.

14. A computer-implemented method of selecting items to recommend, the method comprising:

identifying, for a candidate item, a plurality of reasons for recommending the candidate item to a target user, each reason corresponding to a different respective event or set of events recorded in an event history of the user;

assigning a score to the candidate item responsive to a value probabilistically selected according to a probability distribution, the probability distribution having at least one characteristic that depends on the number of said reasons; and

using the score to determine whether to recommend the item to the target user.

15. The computer-implemented method of Claim 14, wherein the at least one characteristic of the probability distribution comprises entropy.

16. The computer-implemented method of Claim 15, wherein the entropy of the probability distribution decreases as the number of reasons increases.

17. The computer-implemented method of Claim 14, wherein the probability distribution is a normal distribution.

18. The computer-implemented method of Claim 17, wherein the at least one characteristic of the probability distribution comprises a mean configured to depend on the number of reasons, and wherein the score is responsive to the mean.

19. The computer-implemented method of Claim 17, wherein the at least one characteristic of the probability distribution comprises a standard deviation configured to depend on the number of reasons, and wherein the score is responsive to the standard deviation.

20. The computer-implemented method of Claim 14, wherein the score depends upon strength values associated with specific reasons.

21. The computer-implemented method of Claim 14, further comprising assigning a negative feedback score to the candidate item, wherein the negative feedback score is dependent on the number of said reasons.

22. The computer-implemented method of Claim 21, wherein assigning a score to the candidate item further comprises adding the one or more category scores and the negative feedback score.

23. The computer-implemented method of Claim 14, wherein assigning a score to the candidate item responsive to a value probabilistically selected according to a probability distribution further comprises assigning the score such that an overlap of 5 to 20 percent of recommended items between recommendation generation events is achieved.

24. A computer-implemented method, the method comprising:

assigning initial scores to candidate items generated as preliminary recommendations to a target user, the initial scores based on a number of reasons corresponding to different respective events recorded in an event history of the user;

adding probabilistic noise to at least some of the initial scores;

sampling values from a probability distribution that includes the probabilistic noise; and

generating adjusted scores for the at least some of the candidate items, the adjusted scores being responsive to the sampled values.

25. The computer-implemented method of Claim 24, wherein sampling values from a probability distribution that includes the probabilistic noise comprises probabilistically selecting values according to a probability distribution having at least one characteristic that depends on the number of said reasons.

26. The computer-implemented method of Claim 24, wherein the probability distribution is a normal distribution.

27. The computer-implemented method of Claim 26, wherein a mean of the normal distribution depends on the number of reasons, and wherein the score is responsive to the mean.

28. The computer-implemented method of Claim 26, wherein a standard deviation of the normal distribution depends on the number of reasons, and wherein the score is responsive to the standard deviation.

29. The computer-implemented method of Claim 24, wherein generating adjusted scores for at least some of the candidate items further comprises determining a negative feedback score dependent on the number of reasons.

30. The computer-implemented method of Claim 24, wherein generating adjusted scores for the at least some of the candidate items further comprises generating adjusted scores such that an overlap of 5 to 20 percent of recommended items between recommendation generation events is achieved.

What is claimed is:

**1**. A computer-implemented method of selecting items to recommend, the method comprising:

generating a ranked set of candidate items to recommend to a target user based at least partly on item selection actions performed by the target user, the ranked set of candidate items comprising candidate items and associated rankings of the candidate items;

subsequent to said generating, probabilistically varying the rankings of at least some of the candidate items in the ranked set of candidate items automatically by a computer to generate a probabilistically-varied ranked set of candidate items,

said probabilistically varying the rankings comprising using a pseudo-random number generator to vary the rankings by an amount that is based at least in part on a number of the item selection actions performed by the target user;

selecting, from said probabilistically-varied ranked set, a first subset of the candidate items to present to the target user; and

subsequent to said selecting, repeating at least said probabilistically varying the rankings and said selecting to thereby select a second subset of the candidate items to present to the target user, wherein the second subset of the candidate items comprises one or more items that are not members of the first subset of the candidate items.

**2**. The computer-implemented method of claim **1**, wherein selecting, from said probabilistically-varied ranked set of candidate items, the first subset of the candidate items to present to the user comprises selecting a most-highly ranked subset of items.

**3**. The computer-implemented method of claim **1**, wherein at least some of the candidate items are associated with one or more identified reasons for being recommended to the target user.

**4**. The computer-implemented method of claim **3**, wherein the amount which the ranking of a candidate item is probabilistically varied further depends on the number of identified reasons for recommending the candidate item to the target user.

**5**. The computer-implemented method of claim **4**, wherein probabilistically varying the rankings comprises causing candidate items having relatively low numbers of reasons to be probabilistically varied to a greater degree than candidate items having relatively high numbers of reasons.

**6**. The computer-implemented method of claim **1**, wherein probabilistically varying the rankings comprises causing candidate items having relatively low category scores to be probabilistically varied to a greater degree than candidate items having relatively high category scores.

**7**. The computer-implemented method of claim **1**, wherein probabilistically varying the rankings comprises generating values for at least a portion of the candidate items and adding probabilistic noise to the values.

**8**. The computer-implemented method of claim **7**, wherein adding probabilistic noise to the values comprises generating a probability distribution corresponding to each value, wherein the characteristics of the probability distribution depend on the corresponding value.

**9**. The computer-implemented method of claim **8**, further comprising sampling values from each probability distribution to generate a score for said at least a portion of the candidate items, wherein the score is responsive to the sampled value.

**10**. The computer-implemented method of claim **1**, wherein probabilistically varying rankings of at least some of the candidate items in the ranked set further comprises probabilistically varying the rankings to achieve a target overlap value of same recommended items between the first and second subsets.

**11**. The computer-implemented method of claim **1**, wherein said generating, said probabilistically varying, and said selecting are implemented by one or more machines.

**12**. A computer-implemented method of selecting items to recommend, the method comprising:

by a computer system comprising computer hardware:

identifying items selected from an electronic catalog by a target user;

for each of the identified items, selecting associated items from a table of associated items to produce a set of candidate recommendations, said selecting comprising providing an initial score for each candidate recommendation;

for each of the candidate recommendations,

identifying a number of the item selections made by the target user that resulted in the candidate recommendation being selected from the table of associated items, and

calculating a new score for the candidate recommendation using a value selected from a probability function, the probability function configured to provide the value based at least partly on the identified number of item selections, wherein the probability function at least uses a pseudo-random number generator to calculate the new score;

outputting a first subset of the candidate recommendations for presentation to the target user, each of the candidate recommendations being ranked in the output according to the new score for that candidate recommendation; and

subsequent to said outputting, repeating at least said calculating the new score for each of the candidate recommendations and said outputting to thereby output a second subset of the candidate recommendations, wherein the second subset of the candidate recommendations comprises at least some different items from the first subset of the candidate recommendations.

**13**. The computer-implemented method of claim **12**, wherein the at least one characteristic of the probability function comprises entropy.

**14**. The computer-implemented method of claim **13**, wherein the entropy of the probability function decreases as the number of reasons increases.

**15**. The computer-implemented method of claim **12**, wherein the probability function is a normal distribution.

**16**. The computer-implemented method of claim **15**, wherein the probability function is further configured to provide the value based at least partly on a mean configured to depend on the number of item selections.

**17**. The computer-implemented method of claim **15**, wherein the probability function is further configured to provide the value based at least partly on a standard deviation configured to depend on the number of item selections.

**18**. The computer-implemented method of claim **12**, wherein the new scores depend upon strength values associated with specific item selections.

**19**. The computer-implemented method of claim **12**, wherein calculating the new score for the candidate recommendation further comprises assigning negative feedback score to the candidate recommendation, wherein the negative feedback score is dependent on the number of said item selections.

**20**. The computer-implemented method of claim **19**, wherein calculating the new score for the candidate recommendation further comprises determining one or more category scores responsive to the number of item selections in one or more categories selected from the group of a purchase history category, an unbought interests category, and a recent activity category.

**21**. The computer-implemented method of claim **20**, wherein calculating the new score for the candidate recommendation further comprises combining the one or more category scores and the negative feedback score.

22. The computer-implemented method of claim 12, wherein said identifying, said assigning, and said using are implemented automatically by one or more computer processors.

23. A recommendation engine for selecting items to recommend, the recommendation engine comprising:

a computer system comprising one or more processors, said computer system programmed to implement:

a candidate generator configured to:

generate a ranked set of candidate items to recommend to a target user, the ranked set of candidate items comprising candidate items and associated rankings of the candidate items;

a probabilistic scorer configured to, subsequent to the candidate generator generating the ranked set of candidate items, probabilistically vary the rankings of at least some of the candidate items in the ranked set of candidate items to generate a probabilistically-varied ranked set of candidate items, the probabilistic scorer configured to probabilistically vary the rankings by at least using a pseudo-random number generator to vary the rankings by an amount based at least in part on one or more item selection actions performed by the target user; and

a candidate filter configured to select, from said probabilistically-varied ranked set, a first subset of the candidate items to present to the user;

wherein the probabilistic scorer is further configured to probabilistically vary the rankings of at least some of the candidate items a second time to generate a second probabilistically-varied ranked set of candidate items, and wherein the candidate filter is further configured to select a second subset of the second probabilistically-varied ranked set of candidate items to present to the user, wherein the second subset comprises at least some different items from the first subset of the candidate items.

**24**. The recommendation engine of claim **23**, wherein the degree to which the probabilistic scorer probabilistically varies the ranking of a candidate item depends on a number of identified reasons for recommending the candidate item to the target user.

**25**. The recommendation engine of claim **24**, wherein probabilistically varying the rankings comprises causing candidate items having relatively low numbers of reasons to be probabilistically varied to a greater degree than candidate items having relatively high numbers of reasons.

**26**. A method for varying recommendations presented to users, the method comprising:

by a computer system comprising computer hardware:

selecting items as candidates for recommendation to a target user;

ranking the items using a recommendation process that depends at least in part on item selection actions performed by the target user, the rankings configured to estimate an optimal order for presenting the items as recommendations for the target user;

varying the rankings of the items to produce varied rankings by at least using a probabilistic function that depends at least in part on the item selection actions performed by the target user so as to estimate a less optimal order for presenting the items to the target user, the probabilistic function being implemented at least in part by a pseudo-random number generator, such that some of the rankings are increased for some of the items and other of the rankings are decreased for other of the items;

outputting a first subset of the items as a most highly-ranked subset of the items according to the varied rankings as additional recommendations for presentation to the target user; and

subsequent to said outputting, repeating said selecting, ranking, varying, and outputting, based on the same item selection actions performed by the target user, to thereby output a second subset of the items to recommend to the target user, the second subset of items comprising at least some different items from the first subset of items.

**27**. The method of claim **26**, wherein said varying the rankings comprises adjusting the rankings based at least in part on a number of item selections performed by the target user.

**28**. The method of claim **27**, wherein said varying the rankings comprises increasing ranking variability when the number of item selections is relatively higher and decreasing the ranking variability when the number of item selections is relatively lower.