Goal:
To determine the impact of the migration to PrimoVE on the Results/Relevancy Ranking for searches.

Evaluation plan:
1) Develop a set of keywords/search terms that all testing campuses will use
2) First phase: Does the Ranking change with the migration to PrimoVE
    a) Each testing campus will conduct two sets of searches using the provided keywords
        i) One search using PBO
        ii) One search using PrimoVE
    b) Review the first 3 pages of results (a quick literature search mentioned that most users only care about the first 2 pages/20results) - making note of any differences between the two searches
    c) What are the differences between the two sets of search results?
3) Second phase: If significant differences are found between the two instances of primo
    a) Testing Precision for both instances (this is more tricky, since there isn't a predetermined set of what is reverent or not for the results. At least each tester can be consistent with themselves)
        i) For the first 3 pages/30 results, how many are relevant to the search?
            (1) Precision = Number of relevant results/Number of results total
    b) What differences do you note between the precision for both Primos?
4) Final Phase: Qualitative discussion - survey after testing has been completed
    a) What differences did you notice between the results? Did one set of results "feel" more useful than the other?"
    b) What, if anything, seemed confusing in terms of results? What unexpected results did you find?

Andrews Notes:

Metrics:
Look at top twenty results and determine how relevant the results are

Compare them to the original version of Primo

Topical vs User Relevance (prioritising recent results for the user vs relevant to the topic being searched)

Precision vs Recall - recall will be impossible to measure by humans at teh scale of results we retrieve from Primo

5. **Relevance Feedback and Pseudo Relevance Feedback (PSR)**
Here, instead of asking user for feedback on how the search results were, we assume that top k normally retrieved results are relevant. Typical process is as below:
1. Take the results returned by initial query as relevant results (only top k with k being between 10 and 50 in most experiments).
2. Select top 20–30 (indicative number) terms from these documents using for instance tf-idf weights.
3. Do Query Expansion, add these terms to query, and then match the returned documents for this query and finally return the most relevant documents.

https://knowledge.exlibrisgroup.com/Primo/Knowledge_Articles/Primo_Primo_Central_Searching_and_Search_Results_Evaluation

Ranking

- **Relevance**
  - For author and author/title searches, the weight of the original book was increased to ensure that it is ranked higher than material about the book.
  - When searching for author names, the weight of the material written by the author has been increased above the material about the author.
  - Greater weight is given if keywords are found in author, title or subject fields
    - Exact Title match is set to be weighted the most heavily
  - In a keyword search, a record is ranked higher if the term or phrase order is the same as in the record
  - Ranking also factors in how often the search terms occur in the record

- **Academic significance**
  - Published in a peer reviewed journal
  - Number of times cited
  - Material type (i.e. journal article vs. newspaper article)
- **Type of search**
  - If Primo interprets the keywords as a broad topic search, overview material will be higher in the results
  - If Primo interprets the keywords as a known-item search, author and exact title matches will be higher in the results (see "How Do Users Search and Discover" White Paper in Additional Resources below)
  - For citation searches, if a user pastes citation information into the search, Primo's search engine has been enhanced to recognize the author, title, and date better to find the item regardless of the extra metadata in the query.
- **Publication date**
  - More recently published Primo Central materials show up higher in the results list
  - Sites can configure whether or not publication dates affect the relevance of local collections