Appendix to UK Trade Policy: An Independent Assessment

Methodology

1. Introduction

In the UK, the lack of a clearly articulated trade strategy before 2025 meant that the parameters of trade policy—its goals, objectives, priorities, and the tools used to pursue them—had to be inferred. Such information can be gleaned from the government's trade policy actions such as text of free trade agreements, publication of government strategies that intersect with trade, any supporting documentation, and from other relevant policy initiatives. Examining these sources, helps to uncover trends, patterns and overarching themes that emerge from the body of evidence. The presence of a trade strategy is clearly very helpful and potentially central in understanding a government's trade policy, and our report does draw heavily on the trade strategy. However, as our report also makes clear the range and remit of trade policy is remarkably broad, which also means that relying solely on a trade strategy document would be insufficient, and there is an ongoing need for the analysis of a range of documents.

Content analysis is one such method that enables this systematic identification and interpretation of recurring themes. It is used to explain or analyse communicative messages- in this case, in government documentation.[1] The presence of certain words, or concepts mentioned within textual data help identify the focus and intentions of the author or institution. This method is commonly employed across health sciences, political science and public policy disciplines.

2. Data Sources

We refer to policy documents published by the UK governments between 2019-2024 to inform our understanding of the government's trade policy goals, objectives and priorities. These policy documents are published in different forms namely: a Policy Paper, a White Paper and/or a Green Paper. There is a rich literature supporting the use of policy documents as a source to understand policy content and processes, primarily found in health policy and political science disciplines.[2]

---

[1] "Content Analysis Method and Examples" (*Columbia University Mailman School of Public Health*, August 3, 2016) <https://www.publichealth.columbia.edu/research/population-health-methods/content-analysis> accessed November 24, 2025 ; Maier M A, "Content Analysis, Definition Of," *The SAGE Encyclopedia of Communication Research Methods* (SAGE Publications, Inc 2017) <https://doi.org/10.4135/9781483381411.n91> accessed November 24, 2025

[2] Dalglish SL, Khalid H and McMahon SA, "Document Analysis in Health Policy Research: The READ Approach" (2020) 35 Health Policy and Planning 1424 ; Cardno C, "Policy Document Analysis: A Practical Educational Leadership Tool and a Qualitative Research Method" (2019) 24 Educational Administration: Theory and Practice

In the first stage, we identified 33 documents (policy papers and white papers) published by the UK Government between 2019-2024 across different government departments through an online search of government departments, their news bulletins and news items that pointed to the publication of a strategy document. Given the focus of the report on trade policy, we proceeded to identify documents from this list that included discussion of international trade policy. This was done by a close reading of all 33 documents. From this list, we identified 11 documents that discussed the role for trade policy. These 11 documents formed our 'trade corpus'. The documents included in the trade corpus are as follows:

1. Advanced Manufacturing Plan (2023)
2. Net Zero Strategy (2022)
3. Critical Imports and Supply Chains Strategy (2024)
4. UK Export Strategy: Made in the UK, Sold to the World (2021)
5. Levelling Up the United Kingdom (2022)
6. UK Innovation Strategy (2021)
7. National Semiconductor Strategy (2023)
8. UK Critical Minerals Strategy (2022)
9. Integrated Review Refresh 2023 (IR23) (2023)
10. 2025 Border Strategy (2020)
11. International Regulatory Cooperation Strategy (2022)

3. Methodology

In our content analysis, we use a combination of topic modelling and textual analysis—complementary methods that allow us to cross-validate insights and strengthen the robustness of our findings.

3.1 Topic modelling

Topic modelling is a machine learning algorithm which is used to identify semantic patterns or 'topics' portrayed by a text corpus. It processes natural language (text) to identify clusters or groups of semantically similar words or sentences within the body of text. The advantage of topic modelling over other methods such as clustering is the assumption that one document can discuss more than a single topic. It treats each individual document in a collection of texts as a 'bag of words' model. This means that the topic modelling algorithm ignores word order and inherent context and informs its own contextual understanding by observing on how often words occur, and how often they co-occur, within each individual document. Topic models are not synonymous with bag of words, however. While the latter merely counts the presence of words within a collection of documents, the output produced using topic modelling is groups of commonly co-occurring words segregated into sets of topics.

The algorithm to implement topic modelling is written in a programming language - in this case we use Python. It can be executed on the command line or terminal itself or using an Integrated Development Environment (IDE).

The following steps outline the execution of a general topic modelling task, along with some specifications we used for the purpose of this report:

- Step 1: Text preparation

Clean the text to get rid of punctuation, stop words (words such a, a, as, the), trailing white spaces, items that are not text (like symbols and special characters), hyperlinks, standalone numbers (most likely page numbers).

- Step 2: Tokenising and splitting text

Tokenisation refers to breaking the text into smaller units called "tokens." The goal of tokenisation is to split the text into meaningful chunks that can be processed by machine learning algorithms. In this step, the text is split into individual tokens (sentences or groups of words).

- Step 3: Using an embedder and creating embeddings

We use a pre-trained model like BERT (Bidirectional Encoder Representations from Transformers) to create "embeddings" for each token or chunk of text. An embedding is a numerical representation of the token that captures the semantic meaning by representing the token as a vector and the direction of the vector representing its different contextual meanings. For example: if the token 'sea' has 5 different contextual meanings, it will be represented by a vector in a 5 dimensional space. We use BERT's topic modelling package (BERTopic), that uses BERT embeddings to detect topics.

We use the BERT Sentence Transformer instead of others (T5, XLNet, GPT, BART, LLAMA) because BERT uses bidirectional context, meaning it considers both the words before and after a given word to understand its meaning in context. This enables BERT to create much more nuanced representations of text, where semantically similar words (e.g., "automobile" and "car") are mapped closer together, and the same word is represented differently depending on the surrounding words. BERT is also pretrained on vast corpora of text (such as the BooksCorpus and English Wikipedia). This means it has already learned a lot about general language structure, grammar, and meaning. This pretraining allows BERT to understand text at a deeper level, making it a good choice for extracting high-quality features from diverse textual data.

- Step 4: Dimensionality reduction

Once we have generated the embeddings, we need to reduce their complexity to make them easier to visualise and analyse. While these high-dimensional embeddings capture a rich amount of semantic information, they are computationally expensive to process and visualise. Additionally, high-dimensional data can be very sparse and hard to interpret directly.

Dimensionality reduction is the process of reducing the number of features (or dimensions) in the data while preserving the key information. In this case, we use BERTopic's default dimensionality reduction technique UMAP (Uniform Manifold Approximation and Projection).

- Step 5: Clustering

After reducing the dimensionality of the embeddings, we apply clustering algorithms to group similar tokens together into topics. In this process, we use HDBScan (Hierarchical Density-Based Spatial Clustering of Applications with Noise), a clustering algorithm that determines the number of clusters based on the density of the data points. We set the size to 2 to allow each topic to have two similar tokens.

Additionally, we apply semantic filtering based on cosine similarity to further refine the clusters. Cosine similarity measures how similar two vectors are, and we set a similarity threshold of 0.5 to ensure that only tokens that are sufficiently similar to each other are grouped together into the same topic. This step helps create more coherent and meaningful topics from the data.

- Step 6: Representation Model

Once the topics are formed and we have clusters of documents, it's important to represent each topic in a way that is both coherent and informative. To do so we employ a representation model (in this case the Maximal Marginal Relevance Model-MMR). The MMR approach works by balancing relevance and redundancy, ensuring selected documents are both relevant to the topic and diverse from one another. We set MMR to 0.1 which enables us to identify a diverse range of topics, while sacrificing some precision.

We also set a minimum of 2 topics to be detected per document. We do so to capture the cross cutting and multi-dimensional nature of trade policy. This helps us observe (a) how trade interacts with different policy objectives like sustainability, security (b) if different aspects of trade (in the form of policy instruments) are used to address policy objectives. We did not predefine or set a limit on a maximum number of topics, as we adopted an inductive approach in which topics were allowed to emerge directly from the text of the documents in the trade corpus.

The output from this is then stored in a spreadsheet format with columns representing: topic number, topic frequency, top three words in the topic, all the words in the topic, reference sentences/paragraphs from which the topics were created.

3.2 Textual Analysis

Parallel to topic modelling we also undertake a manual and close reading of each of the documents in the trade corpus. We follow an inductive approach to identifying objectives and goals manually. This implies that we allow categories and themes to emerge from the material itself rather than imposing predefined classifications. We first identified narrow objectives, which were then grouped into broader goals based on our expert judgement. This was an iterative process, where we revisited and reanalysed the documents several times and by more than a single researcher, to manage researcher bias. We use qualitative data management software- NVivo to manage, organise and quantitatively represent our findings.

4. Results

Our objective is to be both systematic and to have a procedure which is replicable (other countries, other years, other contexts etc), as well as to allow for nuanced and use expert knowledge to bring a multi-disciplinary focus. Hence, we use both methodologies as they complement each other and allow for verification of results against differing approaches as well as to check for robustness. Depending on our objective: to identify a goal, objective or sector, there are slight variations in the implementation of our methodology.

4.1 Topic Modelling

4.1.1 Identifying Goals

We apply topic modelling on each document, in order to identify which words / terms appear together most frequently and thus 'objectively' seem to have importance. At this stage, we split our document into sentence chunks to identify which groups of sentences are more likely to appear together (and form a topic). The algorithm compares all the sentences across the document it is asked to analyse, and identifies those sentences which are semantically the most similar and then extracts the key words which are most commonly found in those sentences. We have then examined those word groupings and their usage in the text and identified each of the 'key terms' that best encapsulate those groupings. For each document we identify the top 20 key terms. Listed below are examples of the key terms derived from the Advanced Manufacturing Plan and the Critical Imports Strategy.

	Advanced Manufacturing Plan	Critical Imports Strategy
1	R&D	Energy security
2	automotive	national security
3	batteries	skills
4	electricity network	innovation
5	regions / investment zones	academic-business partnerships
6	business environment / costs	trade
7	innovation	infrastructure

8	catapult	climate change
9	aerospace	Trade agreements
10	exports	trade facilitation
11	net-zero / renewable energy	non-tariff barriers
12	automated transport	international coordination
13	funding	economic coercion / geopolitical tensions
14	tax relief	mapping / monitoring
15	critical minerals	finance
16	Digital	diversification
17	free trade	exports
18	resilience	complex supply chains
19	SMEs	investment
20	critical imports	consumer choice, prices, variety

It can readily be seen that doing this provides important information on possible key concerns (energy security, economic coercion); objectives (innovation, trade, regions, resilience), tools (funding, tax relief, trade facilitation) and sectors (automotive, batteries, digital). Of course, many of these terms will have more than one possible interpretation – and that will depend on the context.

4.1.2 Identifying objectives

- For something to be an objective of policy, it has to have a desired outcome. In most cases this is associated with a positive verb, eg: increase growth; boost productivity; enhance competitiveness, and in some cases a negative verb eg: reduce emissions; overcome border barriers; limit movement of illegal goods. Using corpus management and text analysis software- Sketch Engine, we identify all the possible synonyms of such positive and negative verbs. We identify 22 positive and 11 negative verbs.
- We then turn to topic modelling to find words that co-appear with the identified positive and negative verbs. As a 'constrained' approach in comparison to the 'unconstrained' approach while identifying goals, we are only interested in knowing those groups of words that most likely appear with the identified positive and negative verbs. In this instance, we split the text into trigrams or groups of three words, as opposed to a sentence split while identifying objectives. Doing so identifies the following terms across the documents in the trade corpus: exports/new markets, regulations, partnerships/collaborations, hydrogen, barrier removal, investors, competitiveness, decarbonisation, manufacturing, renewable energy, automotive, digitalisation, resilience, technology, economic growth, minerals, and jobs among other terms

The preceding provides an initial list of key objectives as they are related to international trade as derived from the trade corpus. We have then grouped the objectives into broader categories and have also added to the list on the basis of our knowledge and reading of the texts. Our final list of broader objectives (in alphabetical order, as opposed to order of importance), together with their sub-categories is:[3]

1. Distributional issues: Inclusivity / equality
   - Women in business / trade
   - Labour rights
   - Human rights
   - Inequality
   - Fair trade
   - Consumer protection
   - Living standards
   - health
2. Economic growth
   - Innovation and R&D
   - Digitisation
   - Promotion of a given sector
   - Reducing barriers to business and/or trade
     - Simplifying trade at the border
     - Market access
     - Free trade
     - Mobility of workers
     - Rules based trade
   - Infrastructure development
   - Skills, training and education
   - Competitiveness
   - Investment / FDI
   - Wealth creation
   - productivity
3. Economic security
   - National security
   - Security
   - Geopolitical tensions

[3] The items marked in red are objectives that the Topic Modelling approach was unable to identify. Hence, using our judgement, we added additional objectives.

4. Environment / sustainability
    - Sustainable development
    - Energy transition
    - Climate change
    - Circular economy
5. Global cooperation and leadership
    - International development
    - Standards and regulation
    - Science, technology and data?
    - Reputation
6. Resilience
    - Vulnerability
    - Diversification
    - Building domestic capabilities

4.2 Textual Analysis

4.2.1 Identifying objectives

The approach we take here is one of the classic academic literature review method, which involves reading all the documents and doing a summary. The advantage is that it allows the analysis to be more nuanced and draw on the expert knowledge of the analyst.

To assist with organising and analysing the textual data, we used qualitative data management software NVivo. The PDF version of each of the documents from the trade corpus was imported into the software for qualitative coding- a process where you assign labels to segments of textual data.

In identifying objectives manually, we focused on locating those sentences in the text of the document in which the government committed to undertake an action, enact a programme, or expressed a clear aspiration for change. Each of these sentences were then assigned a code or category of objective. For example: The Critical Imports Strategy states, "Government has also pursued mutual recognition agreements (M R A s) with trading partners to facilitate greater market access and ease the importation of goods.", This sentence was assigned the objective codes 'market access' and 'imports', as the sentence itself reflects the aspiration to provide greater markets and make importing easier. This is an example of one sentence reflecting two objectives. The aim was to capture how often a distinct policy objective appeared across the documents, not how frequently a particular verb or phrase was repeated. Therefore, the counting process was based on contextual instances of objectives, not on the number of times specific words appeared. An objective was counted once per context within each document. This means that if the same objective was articulated several times within the same section or discussion, using similar wording or referring to the same intended action, it was treated as a single observation.

We identified 44 distinct objectives across the 11 documents from our trade corpus. These results are very similar to those derived from the topic modelling.

As this was an iterative process, the coded excerpts were reviewed and refined multiple times across multiple coders. Segments that did not clearly express an objective were removed. This iterative refinement ensured consistency and improved the analytical reliability of the coding scheme.

After the contextual coding was completed, NVivo's coding matrix query function was used to generate a systematic count of how often each objective appeared across the documents.

4.2.2 Identifying goals

Using our expert judgement and knowledge of the subject matter, all previously coded objectives were reviewed and compared for similarities in thematic focus. Using NVivo's node structure, objectives were clustered into conceptual groups based on their substantive content. Again, this was an iterative process. The resulting set of goals thus reflects a structured, theme-based interpretation of the government's broader policy intentions. This resulted in the 6 policy goals as mentioned in Chapter 2 of the report.

4.3 Identifying priorities

To identify government priorities from the stated objectives, we needed a factor through which we could compare the relative importance of each of the objectives. This was done by developing a Relative Priority Index (RPI), which enabled the comparison of the contextual frequency of the recorded objectives.

The RPI incorporates both the distribution of each objective (and thereafter goal) across the documents and the overall frequency with which the objective appears. The RPI for each objective was calculated using the following formula:

$$\mathbf{RPI} = \frac{\textbf{Mean of contextual frequency}}{\textbf{Standard Deviation of contextual frequency}} \times \frac{\textbf{Contextual Frequency}}{\textbf{Maximum Contextual Frequency}}$$

The RPI allowed us to compare the scores received for each objective to ascertain which objective/goal was more important than the other. Higher the RPI, greater is the importance of the objective/goal. Important to note that RPI >0. The results of the RPI scores for the goals are discussed in Chapter 2 of the report.

4.4 Identifying sectors

To identify sectors, we rely on a manual approach. We apply a simple tally-mark method: if a sector is mentioned at least once in a document within the trade corpus, we record it as a single count, regardless of how many times it appears elsewhere in that document. Our goal is to capture the breadth of a sector's relevance across different policy areas, rather than the depth of its discussion within a single document. Counting multiple mentions within the same document would introduce bias owing to document length, since longer documents would naturally more likely to contain repeated references. This could distort comparisons across sectors, so we treat each document as a single observation per sector.

We define sectors loosely as a classification of a business activity or industry activity or groups of products where the government has a strategic objective. Examples include: automotive, agriculture, digital etc. In our list of sectors also appear some instances which are conventionally not sectors, but an emerging area of interest for the government. Examples include: critical minerals, critical imports, green industries etc.

5. Conclusion

The results obtained using this methodology are presented in Chapter 2 of the report. During the development of the methodology, we observed the potential for large language models (LLMs) to support textual analysis in the context of trade policy. We experimented with alternative GenAI-powered LLMs, such as ChatGPT, using prompt engineering to provide targeted instructions to identify stated goals, objectives, and priorities. However, compared with the combined methodological approach adopted in this report, these models produced results that lacked sufficient precision and contextual depth.

Read the full report: UK Trade Policy Review: An Independent Assessment:
https://citp.ac.uk/uktpr/uk-trade-policy-independent-review