

Labour Market Intervention Logics in Swedish Activation Programmes

A Dictionary-Based Computational Classification

Louise MacRae

Master of Science in Social Scientific Data Analysis

Master's Thesis

Supervisor: Christopher Swader

Graduate School

May 2026

Abstract

Although there is continuous investment in active labour market policies, the fundamental principles behind publicly funded activation programmes remain poorly mapped. This thesis addresses the gap by developing and applying a theory-driven, dictionary-based computational classification instrument to the full corpus of final project reports from the European Social Fund Programme Area 2 in Sweden, comprising 484 projects and approximately 10,000 pages of documentation. Drawing on established typologies in the active labour market policy literature and the sociology of classification, the study constructs a theoretically grounded typology of intervention logics and classifies each project according to the dominant causal mechanism expressed in its documentation. The results reveal a portfolio characterised by a structured concentration around Work-First and Human Capital logics, with Social Capital occupying the most marginal position and Motivational Activation at a more modest level than its theoretical grounding would predict. The study argues that this distribution reflects institutional selection pressures embedded in the European Social Fund's governance architecture, with the funding system's output-focused accountability infrastructure consistently privileging logics whose causal mechanisms can be translated into administratively legible and quantifiable outputs. Validated against a manual adjudication standard with an effective primary accuracy of 93.3 per cent, the classification instrument demonstrates that theory-driven, dictionary-based approaches can reliably distinguish between distinct intervention logics in a large Swedish-language policy corpus, providing a replicable template for systematic mapping of programme content in contexts where manual analysis is infeasible.

Keywords: Active labour market policy, Intervention logic, Dictionary-based classification, European Social Fund, Computational text analysis, Governance

Word count: 21,492

Popular science summary

Every year, people who are unemployed or far from the labour market receive support through publicly funded programmes. Some get help with language training, others with job placements, rehabilitation, or coaching. But despite billions of SEK being invested in these programmes, we have remarkably little systematic knowledge of the kinds of support being provided and why those approaches are expected to help.

This matters as different programmes rest on fundamentally different assumptions about why people are unemployed. One programme might assume the problem is a lack of skills. Another might assume it is a lack of job contacts, a health-related barrier, or a motivational one. These are not minor differences in approach. They reflect entirely different theories about how people enter the workforce, and result in different interventions. Grouping them under the same label and asking whether they “work” produces little useful knowledge. The answer depends entirely on the type of intervention being evaluated and for whom.

The thesis argues that the portfolio’s composition reflects not just what works, but also what the funding system can see and measure. Programmes whose effects are relational or attitudinal are harder to count and report within the administrative frameworks. These frameworks tend to receive less support as a result. Understanding this matters for anyone who designs, funds, or evaluates labour market policy. The composition of Sweden’s activation portfolio may reflect what the governance system can measure as much as what the research says works. Making that visible is the first step toward asking whether the investment is as well-calibrated as it could be.

Acknowledgements

I would like to thank my supervisor, Christopher Swader, at the Department of Sociology at Lund University, for his guidance and constructive feedback throughout the writing process. I am grateful to the Swedish ESF Council for providing access to the project documentation that forms the empirical basis of this study. I would like to especially thank Petra Ornstein at Ramboll, whose regular engagement with this project, thoughtful feedback, and familiarity with the ESF Programme context have been essential in the construction of this thesis.

Table of Contents

LABOUR MARKET INTERVENTION LOGICS IN SWEDISH ACTIVATION PROGRAMMES	1
1. INTRODUCTION	1
1.1 BACKGROUND	1
1.2 RESEARCH GAP	2
1.3 RESEARCH QUESTIONS.....	3
1.4 DELIMITATIONS	4
1.5 STRUCTURE OF THESIS	4
2. PREVIOUS RESEARCH	5
2.1 ACTIVE LABOUR MARKET POLICY EFFECTIVENESS	5
2.2 TYPOLOGICAL APPROACHES TO LABOUR MARKET INTERVENTION	7
2.3 COMPUTATIONAL TEXT ANALYSIS IN POLICY RESEARCH	8
3. THEORETICAL FRAMEWORK	10
3.1 CLASSIFICATION AS ANALYTICAL AND SOCIO-TECHNICAL PRACTICE	10
3.1.1 <i>Classification as a Tool for Ordering and Knowing</i>	10
3.1.2 <i>Classification as a Site of Stabilisation and Resistance</i>	12
3.2 FROM POLICY THEORY TO INTERVENTION LOGIC	12
3.2.1 <i>Policy as Problem Representation and Causal Assumption</i>	12
3.2.2 <i>Intervention Logic as Theory of Change</i>	13
3.2.3 <i>Mechanisms, Not Activities</i>	13
3.2.4 <i>Functional Orientation and Sequential Activation Trajectories</i>	14
3.3 A THEORETICAL TYPOLOGY OF LABOUR MARKET INTERVENTION LOGICS	16
3.3.1 <i>Human Capital as an Intervention Logic</i>	16
3.3.2 <i>Work-first as an Intervention Logic</i>	18
3.3.3 <i>Capability and Health as an Intervention Logic</i>	19
3.3.4 <i>Social Capital as an Intervention Logic</i>	21
3.3.5 <i>Motivational Activation as an Intervention Logic</i>	23
4. EMPIRICAL SETTING AND DATA	25
4.1 THE EUROPEAN SOCIAL FUND AND PROGRAMME AREA 2	25
4.1.1 <i>Programme Reach and Participant Profile</i>	26
4.1.2 <i>Source and Coverage</i>	27
4.1.3 <i>Document Structure</i>	28
4.1.4 <i>Inclusion Criteria and Exclusions</i>	29
4.2 METHODOLOGICAL JUSTIFICATION OF THE CORPUS	29
4.3 ETHICAL AND ACCESS CONSIDERATIONS.....	31
5. METHODOLOGY	31
5.1 RESEARCH DESIGN	32
5.2 DICTIONARY-BASED CLASSIFICATION	34
5.2.1 <i>Conceptual Derivation of Categories</i>	35
5.2.2 <i>Construction of the Dictionary and Classification Procedure</i>	37
5.2.3 <i>Strengths and Limitations of Dictionary-Based Classification</i>	38
5.3 ITERATIVE ALGORITHM REFINEMENT	39
5.3.1 <i>Initial Observations and Keyword Refinement</i>	39
5.3.2 <i>Differential Weighting of Keyword Indicators</i>	40
5.3.3 <i>Normalisation and Gating</i>	40
5.3.4 <i>Dominance and Mixed Classification</i>	41
5.3.5 <i>Calibration of Secondary Analytical Dimension</i>	41
5.4 LLM-ASSISTED VALIDATION	42
6. RESULTS.....	44

6.1 OVERALL DISTRIBUTION OF PRIMARY INTERVENTION LOGICS.....	45
6.1.1 <i>Primary Classification Results</i>	45
6.1.2 <i>Residual Categories and the Prevalence of Hybrid Logics</i>	46
6.2 HUMAN CAPITAL SUBTYPES.....	47
6.2.1 <i>Human Capital Subtype Results</i>	48
6.2.2 <i>The Dominance of Language Interventions</i>	48
6.2.3 <i>Formal Vocational Education and Validation Interventions</i>	49
6.2.4 <i>Methodological and Substantive Implications of Subtype Concentration</i>	50
6.3 SECONDARY INTERVENTION LOGICS.....	51
6.3.1 <i>Secondary Classification Results</i>	51
6.3.2 <i>Work-First as a Structuring Secondary Logic</i>	52
6.3.3 <i>Secondary Logics and Portfolio Hybridity</i>	53
6.4 FUNCTIONAL ORIENTATION WITHIN THE LABOUR MARKET TRAJECTORY	54
6.4.1 <i>Distribution of Functional Orientations</i>	54
6.4.2 <i>Functional Orientation by Primary Logic</i>	55
6.5 CLASSIFICATION CERTAINTY AND SIGNAL STRENGTH	56
6.5.1 <i>Dominance Ratios by Primary Category</i>	56
6.5.2 <i>Methodological Implications of Dominance Variation</i>	58
7. DISCUSSION	58
7.1 THE INTERVENTION LOGIC COMPOSITION OF PROGRAMME AREA 2.....	59
7.2 THE DISCREPANCY BETWEEN RHETORIC AND CLASSIFIED MECHANISMS	62
7.3 GOVERNANCE, CONCENTRATION, AND THE ARCHITECTURE OF ACTIVATION POLICY	65
7.4 THE RELIABILITY OF THE CLASSIFICATION INSTRUMENT	70
8. CONCLUSION	73
8.1 PRINCIPAL FINDINGS AND THEIR SIGNIFICANCE	73
8.2 WHAT THE VALIDATION EVIDENCE ESTABLISHES	74
8.3 CONTRIBUTIONS	75
8.4 LIMITATIONS	77
8.5 FUTURE RESEARCH	78
8.6 CLOSING REMARKS.....	79
9. BIBLIOGRAPHY	80
10. APPENDICES.....	83
10.1 APPENDIX A	83
10.2 APPENDIX B	88
10.3 APPENDIX C	91

1. Introduction

1.1 Background

Despite decades of sustained investment in Sweden's active labour market policy, uncertainty persists about which types of interventions are being funded (Arbetsförmedlingen, 2026, p. 5). Active labour market policies are state-funded interventions designed to facilitate the transition into employment. This can include job placements, training, subsidised employment, and counselling, and are distinguished from passive income-replacement measures by their orientation toward changing the conditions and capacities that shape labour market participation (Calmfors, 1994). These instruments constitute a central pillar of Swedish social policy, but their internal variety and differential causal assumptions are rarely made explicit (Forslund, 2018).

This opacity is documented empirically at the municipal level by Engdahl, Forslund, and Vikman (2023), who find substantial variation across Swedish municipalities in the types of labour market interventions offered, yet note that these interventions are rarely systematically documented and seldom designed to permit evaluation. Even the broadest available classification of municipal interventions, covering roughly ten categories, proves insufficient to capture the causal mechanisms through which programmes are assumed to work.

The importance of such understandings has been well established. Card, Kluve, and Weber (2018) demonstrate that programme effects vary substantially by intervention type, yet categories through which interventions are classified rarely reflect these meaningful differences in causal mechanism. However, the insight that intervention type matters presupposes that we can identify what type of intervention is being implemented. That presupposition, as this study argues, is rarely examined.

The broader policy context within which these dynamics operate is shaped by the social investment turn in European welfare states, which is a shift toward activation and human capital formation as the primary instruments of social policy (Bonoli & Natali, 2012). The European Social Fund's Programme Area 2 is a direct expression of this logic: it is a financial instrument channelled through national and regional structures to support labour market integration through project-based interventions. Yet the effectiveness of the social investment framework depends on a prior question that has received surprisingly little attention: which interventions are being funded, and through what causal logic? Without a systematic account of programme content, the claim that social investment produces employment gains remains underspecified. It is this gap that the present study addresses.

1.2 Research Gap

Most active labour market policy research concentrates on programme effects rather than programme content, often relying on vague categories that mix interventions with different underlying mechanisms. A training subsidy and a vocational certification programme might fall under the same administrative label while assuming entirely different pathways from activity to outcome.

This gap is directly relevant to the European Social Fund-supported labour market projects in Sweden. While European Social Fund Programme Area 2 constitutes a significant channel for labour market activation, the intervention logics embedded in its project portfolio have not been systematically mapped. As documented in the Swedish context, programmes with distinct logics, differing in their assumptions about how participants become employable, the role of the employer, and what counts as a successful outcome, have routinely been grouped under overhead categories that obscure meaningful variation and make systematic assessment difficult (Forslund, 2018). As a result, it remains unclear whether the European Social Fund's funding is concentrated in specific types of interventions, how

different logics are distributed across the programme, and to what extent projects combine multiple mechanisms within a single intervention.

The aim of this study is therefore to address this gap by developing and applying a theoretically grounded, computationally implemented typology of labour market intervention logics to the European Social Fund Programme Area 2 portfolio.

1.3 Research Questions

The empirical question that follows is how Swedish labour-market activation projects can be characterised in terms of their underlying intervention logics. This is examined in this study through the European Social Fund's Programme Area 2 project portfolio as an empirical case, and the question has both descriptive and analytical value. Descriptively, it produces a structured account of the portfolio's intervention logic composition. Analytically, it generates the basis on which patterns of concentration and marginalisation, meaning whether certain intervention logics account for a disproportionate share of funded projects while others remain structurally underrepresented, can be examined in relation to the governance context that produces them. A second methodological question runs alongside this, concerning whether a dictionary-based, theory-driven instrument can reliably perform this classification at scale in a Swedish-language corpus. Secondary analytical dimensions provide a fuller characterisation of the portfolio's internal structure. This thesis will examine the project's functional role, understood as each project's positional orientation within the labour market trajectory, specifically whether it is designed to stabilise participants prior to activation, prepare them for employment, or facilitate direct labour market entry.

These considerations give rise to the following research questions:

RQ1: How can Swedish labour market activation projects be characterised in terms of their underlying intervention logics?

Sub-RQ: What does this characterisation reveal about the governance context that shapes activation policy?

RQ2: To what extent can a dictionary-based, theory-driven classification instrument reliably distinguish between distinct intervention logics in a large Swedish-language policy corpus?

1.4 Delimitations

The scope of this study is delimited in several ways that influence how the findings should be interpreted. The study is restricted to Programme Area 2 of the European Social Fund for the 2014-2020 programme period. Programme Area 1, which targets skills development among those already in employment, and Programme Area 3, which constitutes a youth employment initiative, fall outside the analytical scope. The study classifies project documentation rather than programme implementation; findings should accordingly be understood as a characterisation of how projects represent their intervention logics in formal documentation, not as a direct account of what participants experienced. The study does not evaluate programme effectiveness or estimate causal effects, and does not link classification outcomes to participant-level data.

1.5 Structure of thesis

The thesis proceeds as follows. Chapter 2 reviews the previous research on active labour market policy effectiveness, typological approaches to labour market intervention, and computational text analysis applied to policy documents, establishing the substantive and methodological gaps the study addresses. Chapter 3 develops the theoretical framework, constructing the typology of intervention logics and situating it within the literature on classification as a socio-technical practice. Chapter 4 provides an overview of the European Social Fund Programme Area 2 programme and its institutional context and describes the corpus of project

documentation used in the analysis. Chapter 5 illustrates the classification methodology, including the construction of the dictionary, the scoring procedure, and the validation approach. Chapter 6 presents the results of the classification across the full project corpus. Chapter 7 discusses the findings in relation to the theoretical framework and the research questions. Chapter 8 presents the final conclusions.

2. Previous Research

The present study sits at the intersection of three bodies of research: the quantitative literature on active labour market policy effectiveness, the smaller but growing literature on typological approaches to labour market intervention, and the methodological literature on computational text analysis applied to policy documents. Each contributes to the analytical foundation of this thesis, and each also reveals a gap that the present study is designed to address.

2.1 Active Labour Market Policy Effectiveness

The most developed body of relevant research concerns programme effectiveness. Card, Kluve, and Weber (2018) find that effects vary substantially by intervention type: placement-oriented programmes produce comparable short and long-term gains, while training-based and private-sector subsidy programmes show stronger returns over the medium to long run, and public-sector subsidised employment tends to produce little or even negative average effects. These differential effects between intervention types are not incidental; they reflect meaningful differences in causal mechanism (Card et al., 2018).

Kluve's (2010) meta-analysis of 137 programme evaluations across 19 European countries reaches broadly consistent conclusions, finding that services and sanctions programmes show a substantially higher probability of estimating a

significant positive employment impact than training programmes. The caveat, however, is that their effectiveness depends critically on the availability of jobs into which participants can be directed (Kluve, 2010, p. 915). Programmes targeting young workers are also consistently less likely to yield positive impacts, suggesting the work-first logic operates differently for populations whose primary barrier is the absence of any prior labour market foothold (Kluve, 2010, p. 912).

In a Swedish context, Forslund (2018) demonstrates that although subsidised employment generally yields positive individual outcomes for participants, significant displacement of unsubsidised workers occurs simultaneously, with displacement tending to be greater the closer the subsidised job is to regular open-market employment. This creates a fundamental tension, as the most beneficial intervention designs for participants may also incur the highest costs on non-participants. Managing this tension requires precise targeting, which in turn depends on a clear understanding of the type of intervention being implemented and its underlying logic.

Together, these studies establish that intervention type matters substantially for programme effectiveness. But they share a critical limitation: the differential effects they document presuppose that we know what type of intervention is being implemented. Classification in this literature typically relies on broad administrative categories that group together programmes operating through fundamentally different mechanisms. The effectiveness literature has not systematically examined whether the categories through which it classifies interventions are adequate for capturing the mechanistic variation that its own findings demonstrate to be consequential. That is the substantive gap this study addresses.

2.2 Typological Approaches to Labour Market Intervention

A directly relevant body of research concerns the classification and typology of labour market interventions. Bonoli (2010) offers one of the most systematic conceptual frameworks, distinguishing between incentive reinforcement measures that strengthen work incentives through conditionality and sanctions, as well as employment assistance measures that reduce obstacles to labour market participation through placement services, job search support, and counselling. This distinction captures something analytically important about how different interventions relate to the labour market, but it operates at a high level of abstraction and does not address how interventions are distributed across these categories in practice within a given programme portfolio, nor does it provide a method for classifying large numbers of interventions systematically.

The Swedish National Audit Office's evaluation of Preparatory and Orienting Education (Riksrevisionen, 2017), a programme procured and managed by the Swedish Public Employment Service, offers a more detailed illustration of what is at stake when administrative categories mistakenly combine distinct interventions. The administrative labels named in the evaluation obscure considerable internal differences in both focus and outcomes (Riksrevisionen, 2017, pp. 5–6). The evaluation demonstrates that, on average, participation in the programme correlates with lower future earnings compared to remaining unemployed or engaging in work practice. However, it should be emphasised that these averages mask significant heterogeneity, with positive effects observed for individuals with low educational attainment or those born outside Europe (Riksrevisionen, 2017, pp. 6–7, 32–33, 36, 38). From a typological standpoint, this underscores two essential analytical dimensions: the type of intervention, meaning the mechanisms activated, and the position within the labour market trajectory (Riksrevisionen, 2017). The evaluation also demonstrates that participants often transition into other programmes as part of a sequential integration process, indicating that interventions should be viewed as steps within chains rather than standalone events (Riksrevisionen, 2017).

What neither Bonoli nor the Swedish National Audit Office's evaluation provides is a scalable method for classifying a large portfolio of interventions according to theoretically derived categories at the level of unique project documentation. Developing such a method is one of the two core contributions of this thesis.

2.3 Computational Text Analysis in Policy Research

The third key area of research relates to computational methods for analysing policy texts. Grimmer and Stewart's (2013) foundational exploration of text-as-data techniques lays the groundwork for the methodological choices in this study. They argue that automated text models should be evaluated by their capacity to perform meaningful social scientific tasks rather than solely by their statistical accuracy. They also claim that dictionary-based methods, when grounded in explicit theoretical foundations and validated against domain-specific standards, can yield substantively interpretable classifications at scales unattainable through manual analysis. A central principle in their work is that all quantitative text models are inherently imperfect due to the context-dependent nature of meaning; however, their validity stems from their ability to generate reliable and theoretically relevant classifications on a large scale. This epistemological stance is directly applicable to the present study, which employs a customised, theory-informed dictionary on a Swedish-language corpus for which no standard instrument exists, thereby placing the responsibility for validation on the researcher rather than on prior validation efforts. Bertoni et al. (2023) expand this framework within the context of computational social science applied to policy analysis, illustrating the variety of scenarios in which text-as-data methods can produce policy-relevant insights from extensive documentary corpora.

Within the computational social science literature, both supervised and unsupervised methods have been used for classifying policy and political texts, and choosing between them has significant implications. Unsupervised topic models,

such as Latent Dirichlet Allocation, uncover hidden thematic structures from a corpus without needing prior definitions, and have been employed on legislative texts, party manifestos, and parliamentary debates to identify dimensions that traditional content analysis might not detect (Grimmer & Stewart, 2013, pp. 5, 15, 18). However, the categories produced by unsupervised models are data-driven rather than based on theory, making alignment with established conceptual frameworks challenging (Bertoni et al., 2023, p. 117). In contrast, dictionary-based and supervised approaches, which require explicitly defining the categories beforehand, are more suitable for research where the theoretical framework is already established and the goal is classification rather than discovery.

The application of computational text analysis on policy documents raises a set of specific challenges that the general methodological literature does not fully cover. Policy texts are primarily produced for administrative purposes, compliance with reporting requirements, and institutional legitimation, rather than for communication between actors. The language they use can differ significantly from the vocabulary of academic or journalistic discourse within which most text analysis tools are developed and validated. Several studies have highlighted this gap. Grimmer and Stewart (2013) point out that applying dictionaries across different domains without revalidation is a source of “serious errors”, and Bertoni et al. (2023) warn that the suitability of any computational approach must be assessed based on the documentary genre being analysed. In this study, this means constructing the keyword dictionary iteratively from the target corpus itself, rather than relying on key general terms from active labour market policy research or English-language policy analysis and validating the classification against domain-specific standards rather than generic benchmarks.

What this methodological literature does not provide is an application of theory-driven dictionary classification to the specific area of labour market intervention logic within a Swedish-language policy corpus. To the author’s knowledge, the present study is the first to do so, and the validation framework it develops in

response to the challenges of this particular domain represents a contribution to the methodological literature.

3. Theoretical framework

Chapter 3 develops the theoretical framework underpinning both the typology and its interpretation. Section 3.1 situates classification as an analytical and socio-technical practice, drawing on Bowker and Star's account of how classificatory infrastructure shapes what is rendered legible as legitimate practices. Section 3.2 establishes the conceptual distinctions between policy as problem representation, intervention logic as a theory of change, mechanisms as distinct from activities, and functional orientation as an independent analytical dimension. Section 3.3 develops the substantive content of the five intervention logics: Human Capital, Work-First, Capability and Health, Social Capital, and Motivational Activation, each grounded in an established theoretical tradition and concluded with an explicit classification criterion.

3.1 Classification as Analytical and Socio-Technical Practice

3.1.1 Classification as a Tool for Ordering and Knowing

Classification is fundamental to knowledge production, as any attempt to analyse complex empirical material requires sorting, grouping, and stabilising categories (Bowker & Star, 1999). Within the sociology of science, classification is not understood as a neutral reflection of reality but rather as a structured way of seeing. Durkheim and Mauss (1903/1963) argue that systems for ordering the natural world reflect social organisation. Douglas (1986) later developed this insight by emphasising that classification systems become stabilised when cognitive conventions align with social institutions. Categories, therefore, persist not because they are objectively true, but because they are institutionally embedded and socially maintained (Bowker & Star, 1999).

In research practice, constructing typologies is an essential act of analytical reduction to render empirical diversity manageable. However, as Bowker and Star argue, classifications do not merely describe the world; they also participate in shaping it (1999). The methodological approach of infrastructural inversion involves shifting analytical attention away from categories as given, and towards the classificatory frameworks that facilitate coordination and make certain interventions legible as legitimate policy responses (Bowker & Star, 1999, p. 34). A written classification typology contains elements of bureaucratic negotiation, compromise, and erasure. Analysing it, therefore, requires attention not only to what is included but also to what is omitted. Granularity, in this sense, is political, since fine-grained systems can enhance visibility for specific groups of mechanisms, whereas coarser systems obscure variation, and residual categories are not neutral but instead shape marginality (Bowker & Star, 1999). When constructing a typology of intervention logics, similar choices are unavoidable, and the framework must remain attentive to how its own structure influences what becomes visible.

This thesis treats classification as an analytical tool designed to identify dominant intervention logics. At the same time, it recognises that such classification inevitably foregrounds certain distinctions while backgrounding others. The study approaches classification not merely as a technical procedure but as an analytical and socio-technical practice, drawing on the sociology of knowledge to examine how categories through which intervention logics are rendered legible are themselves shadowed by the institutional infrastructure within which they operate. This reflexive awareness informs both the construction of the instrument and the interpretation of its outputs and is returned to in the discussion in Chapter 7.

3.1.2 Classification as a Site of Stabilisation and Resistance

When institutions are organised around specific categories, practices may gradually conform to those classifications. If only certain types of activity are measurable or eligible for funding, organisations adapt to fit them. This process leads to stabilisation, as the political struggles that created a classification system fade over time, and categories become accepted without question. When embedded in digital infrastructures, classifications may appear technical rather than normative.

However, convergence is never total. Resistance and reinterpretation occur at the implementation level. Actors negotiate, stretch, or hybridise categories in practice. Even when classifications carry institutional authority, they remain provisional and open to contest (Bowker & Star, 1999).

Recognising this tension between stabilisation and resistance is important when analysing intervention logics, as projects may not fit neatly into assumed ideal types. Instead, these overlaps and hybrid forms should be expected rather than regarded as anomalies, a point that is further elaborated in the interpretation of the Mixed/Hybrid category and secondary classification results discussed in Chapters 6 and 7.

3.2 From Policy Theory to Intervention Logic

3.2.1 Policy as Problem Representation and Causal Assumption

Labour market policies do more than just allocate resources or define activities; they embody structured interpretations of social problems (Astbury & Leeuw, 2010; Pawson & Tilley, 1997). Astbury and Leeuw (2010) describe programmes as ‘theories incarnate’: they carry the expectation that intervention will improve a recurring issue, and they embed assumptions about how and why programme activities will result in change. In this sense, policy acts as applied social theory (Weiss, 1997).

Different paradigms frame unemployment in various ways: as a skills deficit, as weak incentives for job searching, as health-related incapacity, or as a lack of social integration. Each diagnosis implies a corresponding solution (Bonoli, 2010). Labour market programmes, therefore, incorporate causal claims about change, based on the assumption that if the problem is X, then intervention Y will lead to outcome Z. Weiss (1997) makes a clear distinction between implementation theory, which describes how a programme is carried out, and programme theory, which specifies the mechanisms believed to intervene between service delivery and results. It is the latter that the present study aims to reconstruct and classify. This perspective shifts analytical focus from surface-level programme activities to the underlying theories of causation they contain.

3.2.2 Intervention Logic as Theory of Change

The concept of intervention logic offers a framework for revealing these embedded causal assumptions. An intervention logic describes a structured theory of change that links a diagnosed problem to an expected outcome through a specified mechanism (Weiss, 1997; Pawson & Tilley, 1997). Pawson and Tilley (1997) assert that programmes succeed only if they introduce suitable mechanisms to groups within appropriate contexts. Outcomes cannot be fully understood without identifying the generative process that produces them. This means that interventions should be classified not by their actions but by the causal process presumed to generate change, a principle elaborated further in section 3.3.3.

3.2.3 Mechanisms, Not Activities

Administrative classifications of labour market policy typically categorise interventions by activity type: training, work practice, counselling, wage subsidies, or assessment. While useful for budgeting and implementation, activity-based classifications risk obscuring the differences that hold most analytical weight,

namely how change is assumed to occur (Astbury & Leeuw, 2010, p. 365; Card, Kluve & Weber, 2018, p. 896).

The distinction is significant. A programme labelled “training” might aim to develop general human capital, assess existing skills, stabilise daily routines, or motivate disengaged participants. Work practice can serve as a direct employment bridge, a screening tool, or a confidence-building measure. As Astbury and Leeuw (2010) argue, mechanisms appear too frequently as unexplained causal arrows in programme logic models, and collapsing activity and mechanism into a single category is precisely what theory-driven classification should avoid.

Analytical classification, therefore, requires distinguishing activities from mechanisms. Mechanisms refer to the processes through which interventions are expected to generate change, such as changes in skills, incentives, health, social ties, or motivation (Pawson & Tilley, 1997; Astbury & Leeuw, 2010). Organising interventions based on mechanism rather than activity enables comparisons across administrative boundaries and the ability to account for heterogeneous effects, including the differential returns to training, placement, and subsidy programmes documented in the meta-analytic literature (Card et al., 2018). The typology developed in this thesis is organised around this principle: categories are defined by the causal mechanisms believed to generate change, and the classification instrument is designed to identify that mechanism from project documentation rather than simply recognising activity-level vocabulary. The rationale for this design choice is elaborated in Chapter 5.

3.2.4 Functional Orientation and Sequential Activation Trajectories

The distinction between intervention mechanisms and functional position within the labour market trajectory constitutes a second analytical dimension that operates independently of the primary typology. For example, two projects might share the same causal logic, such as Human Capital, yet occupy different roles relative to the

labour market. This distinction is crucial because the same mechanism can operate at different stages of an activation trajectory and involve distinct participant groups, resulting in varied effect profiles even within a single intervention category.

The Swedish National Audit Office's (Riksrevisionen, 2017) evaluation of Preparatory and Orienting Education provides the empirical grounding for this dimension. The evaluation reveals that participants in the programme frequently transition into other interventions as part of a sequential integration logic, meaning that individual projects are best understood not as self-contained events but as steps within chains of activation (Riksrevisionen, 2017). This sequential understanding implies that a project's functional position, specifically whether it is designed to stabilise participants prior to activation, prepare them for labour market entry, or facilitate direct transition into employment, shapes its role within the broader activation architecture independently of the causal mechanism it deploys. The evaluation further demonstrates that average programme effects mask substantial heterogeneity attributable in part to differences in where participants stand in this trajectory at the point of programme entry, with positive effects emerging for the most disadvantaged groups whose primary barrier is distance from, rather than friction within, the labour market (Riksrevisionen, 2017).

Building on this, the present study operationalises functional orientation as a positional dimension capturing each project's role within the labour market trajectory. The categories applied in the classification are derived from the sequential trajectory model implied by the Swedish National Audit Office's evaluation. Stabilising interventions address participants who are not yet ready for activation, preparatory interventions build the competencies or capacities required for labour market engagement, and direct entry interventions are oriented toward immediate employment transitions. A fourth category, assessment, captures interventions whose primary function is the mapping of work capacity or competence rather than facilitating transition toward employment in the conventional sense. Where more than one functional orientation is present, projects

are assigned a mixed/multiple designation. This dimension is not a substitute for the primary logic typology but a complement to it; its analytical value lies in capturing variation in programme design that mechanism-based classification does not alone recover, and its empirical contribution is examined in detail in sections 6.4 and 7.3.

3.3 A Theoretical Typology of Labour Market Intervention Logics

The following five subsections elaborate on the core content of each intervention logic within the typology. Each category is grounded in a well-established theoretical tradition in the active labour market policy literature, defined by the causal mechanism believed to generate change rather than the activities through which the mechanism functions. Each section concludes with a specific classification criterion that outlines the conditions under which a project is assigned to that category in this study. Although the five logics are not mutually exclusive in practice (as noted earlier, hybridity and sequential design are anticipated aspects of activation programming), they are analytically distinct based on their fundamental assumptions about why individuals are unemployed and how that situation can be transformed.

3.3.1 Human Capital as an Intervention Logic

The intervention logic connected to human capital rests on a straightforward premise. Investments in skills boost productivity, which leads to higher wages and more stable employment. Differences in earnings between individuals and countries are mainly seen as differences in the amount of human capital they have acquired (Deming & Silliman, 2025).

A substantial empirical literature has sought to estimate the returns to schooling and other forms of skill acquisition. Across contexts, an extra year of education is

associated with wage increases of about ten per cent, and human capital accounts for a considerable share of wage variation within and between countries (Deming & Silliman, 2025, p. 116). These returns are not fixed or constant. They depend on the equilibrium between supply and demand and are further influenced by technological progress, institutional frameworks, and educational expansion (Deming & Silliman, 2025).

The intervention logic embedded in human capital programmes directly derives from this framework. Individuals invest time and resources in acquiring productive skills: these skills increase their marginal productivity, and higher productivity translates into improved labour market outcomes. The Ben-Porath model formalises this investment logic throughout the life cycle, demonstrating that time spent on skill development lowers short-term earnings but enhances long-term earning potential (Deming & Silliman, 2025). Human capital interventions are thus expected to exhibit a characteristic time profile, with limited or even negative effects on employment in the short term, followed by more substantial impacts over the medium to long run. This pattern is supported by meta-analytic evidence, which indicates that training programmes tend to have modest short-term effects but produce larger impacts two to three years after completion, in contrast to work-first programmes that tend to yield quicker but less enduring returns (Card et al., 2018, p. 928). Human capital programmes appear particularly effective for the long-term unemployed, while work-first approaches tend to deliver stronger results for more disadvantaged groups, implying that effectiveness relies on the alignment between intervention logic and participant characteristics (Card et al., 2018).

In the deductively derived typology developed in this thesis, a project is classified as Human Capital when its primary theory of change centres on skill formation as the mechanism through which employability and earnings are expected to improve. The defining feature is not simply the presence of training activities, but the underlying assumption that investment in productive capacities is the central

pathway to labour market integration, whether through formal qualification, language acquisition, validation of prior learning, or structured vocational training.

3.3.2 Work-first as an Intervention Logic

The work-first intervention logic represents a fundamentally different theory of change compared to the human capital framework. While human capital programmes mainly focus on investing in productive capacity, work-first approaches centre on labour market attachment as the primary objective, addressing barriers, incentives, and informational frictions that prevent unemployed individuals from gaining employment. The core assumption is not that participants lack skills, but that they lack a connection to the labour market. Reconnecting them through placement services, job search support, counselling, sanctions, and incentive adjustments is seen as the most direct route to employment.

Bonoli (2010) provides a conceptual vocabulary that helps clarify the internal structure of this logic. He identifies two distinct but related mechanisms. The first, which he calls incentive reinforcement, includes measures that strengthen work incentives for benefit recipients through conditionality, time limits on reciprocity, and sanctions. The second, employment assistance, involves measures that reduce obstacles to labour market participation, such as placement services, job-search programmes, counselling, and job subsidies that increase the chances of an unemployed person making contact with a potential employer (Bonoli, 2010). While these mechanisms are analytically separate, they tend to operate as complements; incentive reinforcement encourages individuals to enter the labour market, while employment assistance reduces the barriers that would otherwise hinder successful transitions.

The work-first logic rests on a theory of unemployment that attributes its cause not mainly to skill deficits, but to the structural and behavioural conditions that prevent labour market attachment. Misaligned incentives and search frictions are prominent

among these factors. From this perspective, passive income replacement risks deepening labour market detachment by reducing the expected returns to job searching. Kluve (2010) offers empirical evidence for this concern. His meta-analysis finds that ‘Services and Sanctions’ programmes (which broadly align with work-first methods) demonstrate a considerably higher likelihood of showing a significant positive impact on employment than training programmes, both in the overall sample and in the post-1990 evaluations (Kluve, 2010, pp. 911–912).

The effectiveness of work-first programmes is not unconditional, however. Kluve finds that their success depends on the availability of jobs into which participants can be directed (2010, p. 915). This points to a fundamental demand-side constraint. Where obstacles to employment are primarily structural, work-first approaches risk pushing individuals toward a labour market that cannot absorb them, producing displacement rather than durable integration. Programmes targeting young workers are also consistently less likely to yield positive impacts (Kluve, 2010, p. 912), suggesting the work-first logic operates differently for populations without an established labour market presence.

In the typology developed in this thesis, a project is classified as Work-First when its primary theory of change centres on rapid labour market attachment through workplace exposure, with internship, wage subsidy, or direct job entry constituting the central mechanism linking programme activities to employment outcomes. The defining feature is not the absence of any skill-related component, but the underlying assumption that the central obstacle to employment is a structural barrier to labour market attachment, and that direct exposure to the labour market is the primary causal pathway through which outcomes are expected to improve.

3.3.3 Capability and Health as an Intervention Logic

Capability-oriented approaches conceptualise labour market disadvantage primarily as a restriction on individuals’ real freedom to secure employment.

Building on Sen's capability approach, the emphasis shifts to what people can actually do and become, highlighting their substantive opportunities to achieve valued ways of being and doing (Bonvin & Farvaque, 2005, p. 270). From this standpoint, access to resources such as training or employment is necessary but insufficient, since individuals vary greatly in their ability to transform these resources into actual labour market participation. The key question is not simply what opportunities a person has, but whether they can genuinely utilise them.

This introduces a distinct causal mechanism centred on conversion factors. Individual characteristics (e.g., health status, disability, or psychological well-being), as well as social and institutional conditions, shape the extent to which a person can translate available resources into effective labour market participation. As Bonvin and Farvaque (2005) argue, meaningfully assessing a job-seeker's situation requires evaluating their actual capability to work and identifying what institutional action is needed to reduce the distance between their current circumstances and the labour market.

For many individuals, labour market disadvantage is rooted in health-related or functional constraints compounded by structural barriers that limit their capacity to engage in employment. Evidence on persons with disabilities illustrates this clearly. The employment rates for this group average nearly 30 per cent below those of non-disabled persons, and illness or disability is the most cited reason for inactivity among working-age persons with disabilities (Sainsbury et al., 2017, pp. 93–94). The corresponding intervention logic focuses on enhancing the individual's capability to work, rather than solely building skills or securing job placements. This involves both strengthening individual functioning and modifying the conditions in which participation takes place, through vocational rehabilitation, medical or psychological support integrated with employment services, job coaching, workplace adaptations, and personalised support schemes (Sainsbury et al., 2017).

A key implication of the capability approach is that interventions must be context-sensitive and individually tailored, since capability constraints vary across individuals and cannot be adequately addressed through standardised programmes. Bonvin and Farvaque (2005) draw attention to the “micro turn” in social policy, a shift towards individualised, tailor-made measures in which the appropriate intervention depends on an accurate, situated assessment of the person’s capability set, itself shaped by the informational basis available to local actors.

In the typology developed in this thesis, a project is classified as Capability and Health Activation when its primary theory of change centres on reducing health-related or functional barriers to labour market participation, with rehabilitation or assessment of work capacity constituting the central mechanism linking programme activities to employment outcomes. The defining feature is that the intervention operates on the conditions that enable individuals to convert opportunities into participation, rather than primarily on the opportunities themselves or on the skills required to make use of them.

3.3.4 Social Capital as an Intervention Logic

Social capital approaches understand labour market disadvantage as a structural problem of access, specifically the absence of social networks and the resources embedded within them. Putnam defines social capital as the connections among individuals that encompass both social networks and the norms of reciprocity and trustworthiness that arise from them (Putnam, 2000). What matters, to this account, is not only what individuals know or can do, but who they know and what those relationships make possible. Social capital in this sense is a collective resource, residing in relationships and communities rather than individuals (Putnam, 2000).

The core mechanism through which social capital shapes labour market outcomes is the structure and function of social networks. Putnam (2000) distinguishes bonding social capital, characterised by dense, inward-looking networks among

homogenous groups, and bridging social capital, which refers to more outward-facing connections across socially distinct groups. This distinction maps closely onto Granovetter's (1983) framework of weak and strong ties. Strong ties, maintained through frequent and emotionally close contact, generate cohesive clusters of relationships. Weak ties, found between acquaintances rather than close friends, bridge otherwise disconnected segments of the social system and carry non-redundant information that strong ties typically cannot provide (Granovetter, 1983).

It is this bridging function that is of particular significance for labour market access. Individuals whose networks consist primarily of strong ties are effectively confined to the information circulating within their immediate social world, placing them at a disadvantage in contexts where finding employment depends on knowing about the right opportunities at the right time. Granovetter (1983, pp. 205–206) illustrates that employees are significantly more likely to learn about new employment opportunities via weak ties rather than strong ones. Furthermore, the use of weak ties correlates positively with higher occupational achievement. Conversely, the substantial aggregation of social energy within strong ties serves to fragment communities into insular networks, and such isolation constitutes one of the factors contributing to the persistent nature of poverty.

This logic defines a coherent intervention rationale. Rather than building individual productive capacity through training or securing immediate employment through work-first activation, social capital interventions aim to expand and restructure participants' network positions to improve access to opportunities. In practice, this may involve facilitating contacts with employers and professional networks, building bridging ties across social groups, or connecting disadvantaged individuals to communities of practice where labour market information and resources circulate. The intervention, in this logic, targets the relational architecture that mediates access to opportunity. It should be noted, however, that not all networks generate bridging capital; tightly bonded networks may reinforce in-group solidarity without translating into broader labour market access (Putnam, 2000).

For this reason, social capital mechanisms tend to function as enabling conditions within more complex intervention designs rather than complete standalone theories of change.

For the purposes of this typology, a project is classified as Social Capital when its primary theory of change is the expansion of participants' social networks, the facilitation of bridging ties, or the structured connection of disadvantaged individuals to network-embedded resources and opportunities. The defining feature is that the intervention targets the relational architecture mediating access to opportunity rather than building individual capacity or facilitating direct labour market attachment.

3.3.5 Motivational Activation as an Intervention Logic

Motivational and behavioural activation approaches conceptualise labour market disadvantage as a problem of decision-making, attention, and choice architecture. Where human capital and social capital logics target the objective conditions structuring labour market access, behavioural activation targets the psychological and cognitive processes through which individuals perceive and respond to labour market incentives and opportunities. The defining premise is that individuals do not always act as fully rational, optimising agents. Systematic deviations from rational choice arising from cognitive biases, present bias, inertia, and inattention can contribute to suboptimal labour market outcomes independently of individuals' objective endowments or constraints (Chetty, 2015).

The foundational framework for this logic is the behavioural economic distinction between decision utility and experienced utility. As formalised by Chetty (2015, pp. 4–6), the objective an agent maximises when making choices may diverge from the well-being they experience as a function of outcomes. This divergence, which may arise from present bias, inattention, or cognitive limitations, generates what the behavioural public economics literature terms an internality. This is a harm

individuals impose on themselves through systematically suboptimal decision-making. The existence of internalities provides a distinct rationale for policy intervention, one that is conceptually separable both from classical externality arguments and from the assumptions underlying standard neoclassical activation models.

Sunstein and Thaler's (2003) framework of libertarian paternalism, integrated into a broader policy-analytic framework by Chetty (2015), holds that choice architecture (the design of the environment in which decisions are made) systematically shapes outcomes without requiring altered incentives or restrictions on freedom of choice. The most influential instrument within this framework is the default rule. By changing what happens when an individual takes no active decision, policymakers can substantially shift outcomes among those who exhibit inertia, without affecting the choices of those who actively deliberate. Applied to the labour market, this logic motivates interventions such as automatic enrolment in activation programmes, simplified information about available support, and the removal of administrative complexity that discourages take-up of entitlements. In practice, motivational and behavioural mechanisms commonly function as enabling conditions that precede and facilitate other logics.

In the typology developed in this thesis, a project is classified as Motivational Activation when its primary theory of change targets cognitive or motivational barriers to labour market engagement. This includes the reduction of complexity and administrative friction, simplified information provision, the restructuring of default enrolment rules, and motivational support. The rationale for including this category rests on the accumulating evidence that behavioural barriers constitute an independent and substantively important source of labour market disadvantage, and that addressing them requires a qualitatively distinct set of policy instruments.

4. Empirical Setting and Data

Chapter 4 describes the empirical setting and data. Section 4.1 provides an overview of the European Social Fund and Programme Area 2, including its programme reach and participant profile. Sections 4.2 through 4.4 describe the final report corpus, covering source coverage, document structure, inclusion criteria, methodological justification, and ethical considerations.

4.1 The European Social Fund and Programme Area 2

The European Social Fund serves as the principal financial instrument of the European Union dedicated to promoting employment, social inclusion, and the development of human capital across its member states. In Sweden, administration of the European Social Fund is managed by the Council for the European Social Fund in Sweden (Swedish ESF Council), a governmental agency operating under the Ministry of Employment. This body is tasked with issuing calls for applications, distributing funding, and overseeing the processes of follow-up and evaluation (Svenska ESF-rådet, n.d.). The fund allocates resources within each programme period to projects implemented at regional and national levels, co-financing initiatives that are required to align with both the strategic frameworks established at the European Union level and the policy priorities pertaining to the national labour market (Svenska ESF-rådet, 2022).

The programme period studied in this thesis, 2014-2020, was organised around three programme areas. Programme Area 1 targeted skills development for those already in employment, focusing on organisational competence and the relationship between education and working life. Programme Area 3 constitutes a youth employment initiative directed at regions with high youth unemployment rates. Programme Area 2, which forms the empirical scope of this study, was oriented toward increasing transitions to employment among individuals who are far from the labour market (Svenska ESF-rådet, 2022). Programme Area 2 projects were

permitted to pursue either direct employment outcomes or the development of methods to facilitate transitions toward work, reflecting a deliberate flexibility in the programme's theory of change.

The overarching aim of Programme Area 2 was to support individuals who stand far from the labour market, and to develop and apply methods that facilitate their transitions to employment or, at a minimum, movement toward the labour market. These projects were not required to achieve direct employment outcomes, and could instead focus on methodology development and structural capacity-building. Common project owners included the Swedish Public Employment Service, municipalities, and civil society organisations, with municipalities particularly prominent given their statutory responsibility for social assistance recipients, many of whom formed the programme's core target group. The Public Employment Service and municipalities thus operated as both project owners and cooperation partners throughout the portfolio, a structural feature that shaped both the scope and the framing of funded interventions (Sweco, 2021).

4.1.1 Programme Reach and Participant Profile

The empirical scale of Programme Area 2 over the 2014-2020 period was significant. A total of 210,629 individuals participated in Programme Area 2 projects during the programming period, where three target groups were formally prioritised and together made up the majority of participants: young people aged 15-24, newly arrived immigrants, and recipients of economic assistance from municipalities (Ramboll, 2025). These categories were not mutually exclusive, with considerable overlap across groups. The participant flow over time reflected wider policy dynamics. Youth were the dominant group in the early years of the programme, while the proportion of newly arrived immigrants increased markedly during 2017-2019, coinciding with the period following the 2015 increase in refugee arrivals.

The projects in Programme Area 2 varied substantially in scope. The median project enrolled approximately 97 participants, but the distribution was heavily right-skewed, with some projects exceeding a thousand participants and 38 projects enrolling fewer than ten. In terms of primary target group composition, 53 projects were oriented primarily toward newly arrived immigrants, 100 toward recipients of economic assistance, and 90 toward young people. There was substantial overlap between groups: 53 projects targeted combinations of newly arrived immigrants and economic assistance recipients. A substantial group of 83 projects could not be classified under any single dominant target group, of which 16 were primarily oriented toward recipients of activity support (Ramboll, 2025). This compositional heterogeneity in target group orientation provides important context for interpreting the portfolio's intervention logic distribution, even if the present study does not systematically link participant composition to classification outcomes at the project level.

4.1.2 Source and Coverage

The primary data source for this study consists of project description documents drawn from the complete Programme Area 2 project portfolio for the 2014-2020 programme period. These documents were obtained directly from the Swedish ESF Council through a contact at the agency and take the form of final project reports, each submitted by the project owner upon completion. The initial corpus comprised 517 files corresponding to the full set of funded Programme Area 2 projects. Following pre-processing, 33 documents were excluded on technical grounds, as described in section 4.1.4. The final analytical corpus consists of 484 documents representing the same number of projects.

Although the final reports are not freely accessible via a public database or official archive, they were produced using public funds within a publicly administered programme and are not classified or subject to secrecy provisions. Their restricted

accessibility reflects administrative rather than legal constraints, and their use here is consistent with the public accountability purpose for which they were created.

4.1.3 Document Structure

Each final project report includes several structured sections: a project summary, a description of the project idea and anticipated outcomes, an account of actual results and achievements in relation to predefined indicators, a detailed account of methods and working practices, as well as sections dedicated to communication, dissemination, follow-up, and evaluation. Although the document template varies across the portfolio, it adheres to a broadly consistent structural framework, thereby facilitating comparison.

For the purposes of this study, the entire text of each document was used as the unit of analysis rather than individual sections. This decision was made on the grounds that intervention logic is expressed diffusely throughout project documentation and cannot be reliably confined to any singular section. While the project description and methods sections tend to contain the most mechanism-specific language, indicators, result accounts, and summary paragraphs also encompass terminology directly pertinent to the classification of the underlying causal assumptions of an intervention. Employing the complete document maximises signal coverage and obviates the need for secondary decisions regarding which sections to include or exclude, whereas the minimum length threshold mitigates noise from documents whose textual content is insufficient to support classification.

The metadata variables directly available from the project documents include, at minimum, project title, project period, project owner, and the full text of each structured section. Additional metadata, including regional location, funding amount, and registered target group, are available at the programme level but were not incorporated in the present analysis, which focuses exclusively on the textual content of the reports as the basis for classification.

4.1.4 Inclusion Criteria and Exclusions

Inclusion in the final corpus required that a document could be successfully extracted and that the resulting plain text contained at least 500 words or 1000 characters. This threshold was set to ensure that the classification instrument had sufficient textual signal to operate reliably. Documents falling below this floor are unlikely to contain enough mechanism-specific terminology to produce a meaningful intervention logic score. The excluded documents (n=33) did not exhibit any obvious systematic bias in terms of project type or period, and their exclusion is not expected to introduce material distortion into the distributional results. The final corpus of 484 documents represents 93.6 per cent of the original portfolio.

4.2 Methodological Justification of the Corpus

The use of standardised project reports as the analytical corpus is methodologically well-suited to the research questions for several reasons. First, project documentation is the primary medium through which funded projects articulate their own rationale, which makes it a direct window onto how intervention logics are constructed and expressed. Project owners are required to describe not only what activities they implemented but why those activities were expected to produce change, a requirement that generates precisely the kind of mechanism-oriented language that a theory-driven classification instrument is designed to detect.

Second, the standardisation of the report template across the Programme Area 2 portfolio creates a degree of comparability that would be difficult to achieve with a more heterogeneous corpus. While the depth and quality of individual reports vary, the shared structure means that certain types of content appear in predictable locations across documents, facilitating consistent extraction and reducing idiosyncratic variation in what the text covers. The standardised template, however, is not a neutral container. As Bowker and Star (1999) argue, reporting infrastructures shape what can be expressed within them as much as they facilitate

comparison across cases. The shared structure of the final report template reflects prior decisions about what counts as relevant activity and how projects are expected to narrate their own rationale, which were made independently of the interventions they document. The corpus is therefore not simply a window onto intervention logics as they exist in practice but a record of how those logics were rendered expressible within a specific administrative framework, a constraint whose implications for what the classification can and cannot recover are addressed in the discussion in Chapter 7.

Third, the volume of the corpus is large enough to support meaningful distributional analysis at the portfolio level, while remaining within the range that computational text analysis methods can handle reliably. Manual coding at this scale would be prohibitively resource-intensive and likely to introduce unacceptable levels of coder drift, making the computational approach not merely convenient but analytically necessary.

A relevant limitation of this data source concerns the relationship between what projects describe and what they actually do. Final project reports are submitted by project owners and may be subject to retrospective framing, strategic self-presentation, or alignment with programme requirements that do not fully reflect the interventions as implemented on the ground. The key analytical implication is that the classification reflects the textually expressed intervention logic, being the causal assumptions embedded in how projects describe themselves rather than a direct observation of programme content. Whether this textual logic accurately maps onto implementation is a question the present study cannot resolve, but the textual expression of intervention logic is itself a substantively meaningful object of analysis.

4.3 Ethical and Access Considerations

The data used in this study consists of administrative documents produced within a publicly funded programme and is not personal data. Project reports do not contain information about individual participants; they describe aggregate project activities, outcomes, and methods at the project level. The names of project contact persons appear in the reports as formal administrative information, but are not used in the analysis and are not reproduced in the thesis. This study involves neither human subjects nor sensitive personal data.

The documents were obtained directly from the Swedish ESF Council, with the contact facilitated through Ramboll, which acts as an external evaluator for Programme Area 2. The documents were shared specifically for the purpose of developing a classification of intervention logics within the Programme Area 2 portfolio, and no conditions were placed on research use beyond this intended purpose. The classification developed in this thesis is intended to inform future evaluation work across multiple European Social Fund programme areas, enabling evaluators to assess effectiveness and results across theoretically derived categories of intervention, though this application lies outside the scope of the present study. The thesis itself does not evaluate European Social Fund Programme Area 2 or assess programme effectiveness; the project reports are used solely as a corpus for developing and applying a typology of labour market intervention logics. The documents are not classified, and no confidentiality obligations attach to their content. No specific project reports are referenced or quoted in this thesis.

5. Methodology

Chapter 5 describes the methodology. Section 5.1 sets out the research design. Section 5.2 introduces dictionary-based classification and its conceptual foundations. Section 5.3 documents the iterative refinement of the algorithm, including keyword weighting, gating, dominance rules, and the calibration of the

functional orientation dimension. Section 5.4 presents the LLM-assisted validation procedure and results.

5.1 Research design

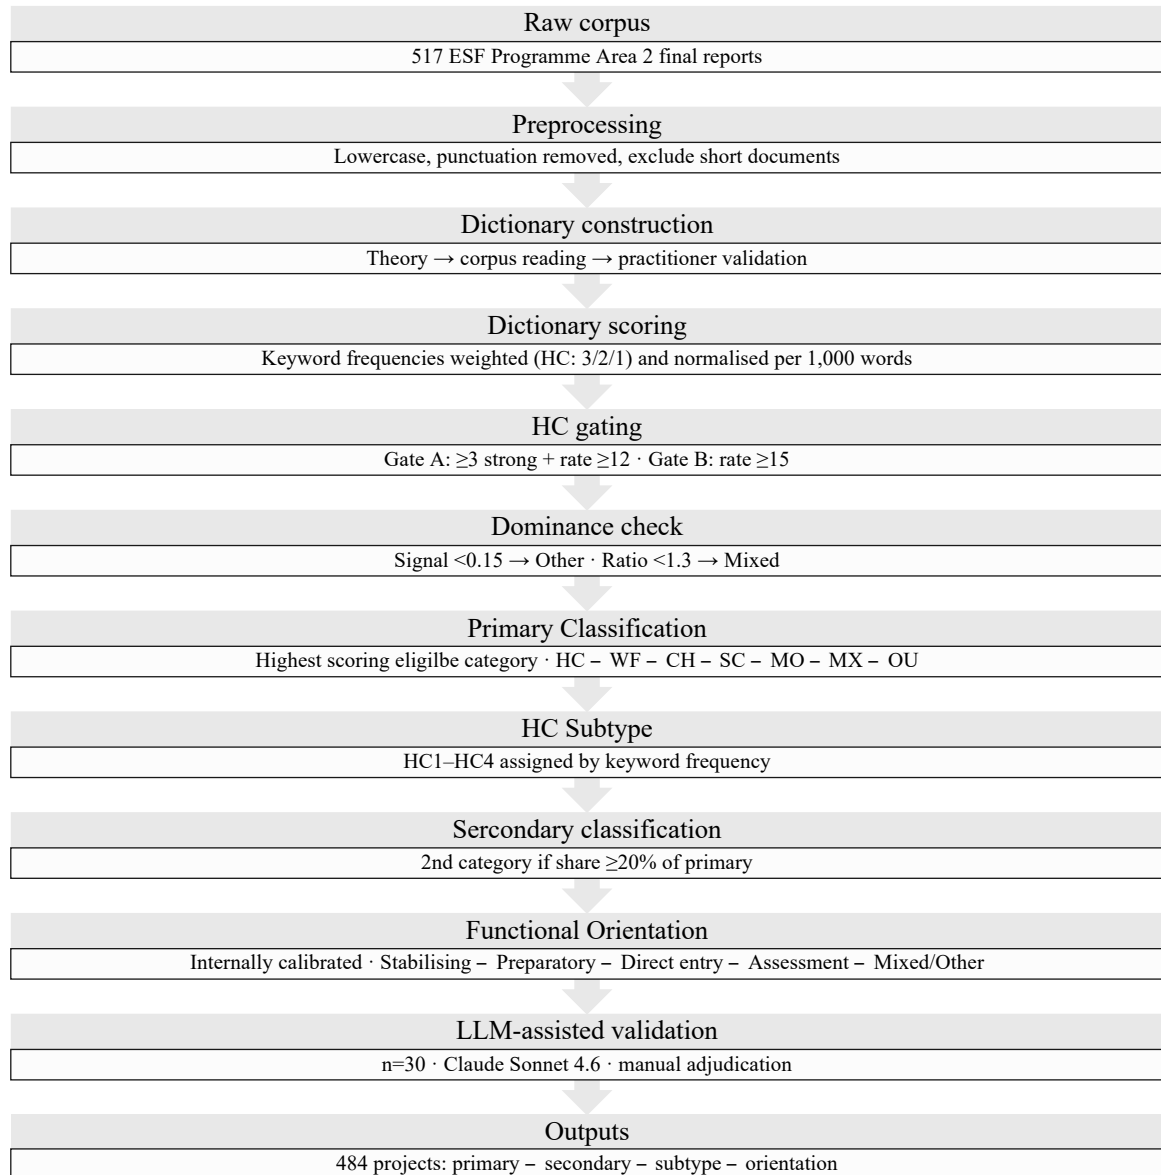
This thesis adopts a theoretically informed, dictionary-based text classification approach to derive a structured typology of intervention logics within projects funded under the European Social Fund Programme Area 2. The study is positioned at the intersection of active labour market policy research and computational text analysis, combining conceptual typology-building with systematic document classification. Within the broader field of Computational Social Science (CSS), this corresponds to a “text as data” design in which automated methods are used to summarise and structure large corpora of texts for substantive social science questions, rather than to model language in full detail (Bertoni et al., 2023; Grimmer & Stewart, 2013).

The primary objective of the research is to identify and categorise the dominant intervention logics embedded within projects funded by the European Social Fund. Rather than evaluating programme outcomes or estimating causal effects, the analysis focuses on the internal policy logic expressed in project documentation, treating automated text classification as a tool for organising project texts into theoretically meaningful categories (Grimmer & Stewart, 2013).

The unit of analysis is individual Swedish European Social Fund Programme Area 2 project descriptions extracted from final reports. These project texts constitute formalised representations of funded interventions and provide insight into how labour market activation strategies are operationalised at the implementation level. The dataset consists of 484 projects drawn from 517 source files, comprising 10 037 pages of project documentation and approximately 4.3 million words following pre-processing exclusions.

Keywords associated with distinct intervention logics (See Table A1 in Appendix A) were constructed through deductive derivation from the theoretical framework established in Chapter 3 and subsequently expanded through iterative reading of project descriptions to identify context-specific Swedish terminology. The keywords and intervention logics were also refined through consultation with practitioners at Ramboll and the Swedish European Social Fund Council, drawing on their familiarity with the corpus and the programmatic language used across the Programme Area 2 portfolio. Projects are classified according to the relative dominance of these theoretically informed dimensions. Throughout this thesis, the term instrument refers to this complete classification system, while classification algorithm refers specifically to the computational procedure implemented in R. Figure 1 provides an overview of the complete classification pipeline, from corpus preparation through to final output.

Figure 1: Overview of the Dictionary-Based Classification Pipeline



5.2 Dictionary-based classification

Dictionary-based classification forms part of a broader set of computational text analysis techniques within CSS, particularly in contexts where categories are theory-driven and require interpretability (Bertoni et al., 2023; Grimmer & Stewart, 2013). Rather than allowing patterns to emerge inductively through machine learning, the classification framework is constructed deductively from established labour market intervention theory (Grimmer & Stewart, 2013). This corresponds to

a distinction in data analysis between summarisation approaches, which are aimed at identifying patterns and structures in data, and predictive or causal approaches, which prioritise forecasting or inference (Bertoni et al., 2023; Grimmer & Stewart, 2013). This design aligns with the principle that text models should be evaluated based on their ability to perform a useful social scientific task. Here, that entails assigning documents into theoretically relevant categories rather than estimating treatment effects (Grimmer & Stewart, 2013).

5.2.1 Conceptual Derivation of Categories

The dictionary was developed based on the intervention dimensions identified in Chapter 3. Drawing on labour market intervention theory and typological research in the active labour market policy literature, interventions were conceptualised along the following primary dimensions (see Table 1).

Each dimension represents a distinct intervention logic, defined by its causal mechanism, rather than a specific activity type. The focus is, therefore, not on surface-level programme components such as “training”, “mentoring”, or “subsidy”, but rather on the fundamental mechanism implied by the intervention. This theoretical foundation ensures that the classification captures the underlying policy logic rather than merely administrative labels, aligning with the focus on designing models that facilitate the measurement of theoretically pertinent concepts instead of relying solely on model fit or predictive content (Grimmer & Stewart, 2013). This methodology reflects the necessity for problem-specific validation to ensure that word lists accurately mirror language use within the particular policy domain.

Table 1: Theoretical Typology of Labour Market Intervention Logics

Type	Label/Code	Subtype	Mechanism	Inclusion criteria
Human Capital	HC1	Credential / Certification	Increase formal qualifications through certification or professional licensing	Certification, diploma or formal qualification explicitly stated as outcome
	HC2	Language Training	Increase employability via language acquisition	SFI or vocational language central to intervention
	HC3	Validation of Prior Learning	Recognise and formalise existing competence	Validation or real competence is core activity
	HC4	Formal / Vocational Education	Structured curriculum-based vocational programme	Longer structured training with defined curriculum
	HC5	Other Structured Training	Structured training without formal credential	Curriculum-based training central but no certification
Work-First	WF	NA	Rapid labour-market attachment through workplace exposure	Internship, wage subsidy or direct job entry central
Capability and Health	CH	NA	Reduce health-related or functional barriers	Rehabilitation or work capacity assessment central
Social Capital	SC	NA	Access employment via relational capital	Networks expansion or bridging ties central
Motivational Activation	MO	NA	Change behavioural readiness or attitudes	Coaching or motivational barrier reduction central
Mixed / Hybrid	MX	NA	Two logics with similar strength	Dominance ratio threshold not met
Other / Unclear	OU	NA	No clear dominant logic	Low signal density across categories

5.2.2 Construction of the Dictionary and Classification Procedure

For each intervention dimension, a corresponding keyword dictionary was constructed by first deriving an initial list of terms directly from the theoretical definitions and inclusion criteria established in Chapter 3, then expanding it through iterative reading of project descriptions across the portfolio to capture the Swedish-language vocabulary through which each logic is expressed in practice. The final step involved consulting with practitioners at Ramboll and the Swedish ESF Council to validate that the terms reflected programmatic language as used across the portfolio (Grimmer & Stewart, 2013). This method uses the relative frequency of keywords to measure the presence of specific categories within the project descriptions. The dictionary includes core conceptual terms (e.g., “skills development”, “capacity building”), operational indicators (e.g., “training modules”, “certification”), as well as synonyms and semantically related expressions identified through iterative reading of project descriptions.

The construction process followed an iterative human-in-the-loop workflow, where automated methods were used to augment and amplify the researcher’s close reading of the text (Grimmer & Stewart, 2013). Initial keyword lists were derived from theory and refined through exploratory reading to identify context-specific terminology. This bespoke construction is a direct response to Grimmer and Stewart’s warning that “serious errors” occur when dictionaries are imported from unrelated domains without accounting for domain-specific language use (Grimmer & Stewart, 2013, pp. 8–9).

Care was taken to distinguish between keywords signalling primary intervention logic and those referring to secondary activities and contextual rhetoric. For example, the presence of education terminology does not automatically indicate a human capital logic unless the project framing explicitly positions skill acquisition as the primary mechanism of change (Grimmer & Stewart, 2013). Each project description was processed through a classification algorithm by calculating the relative frequency and salience of keywords associated with each intervention

dimension. In accordance with Grimmer and Stewart's treatment of text as data methods as task-oriented and interpretability-enhancing tools, projects were assigned a primary intervention logic based on the highest scoring dimension, and, where relevant, a secondary logic where the second-highest scoring dimension exceeded a predefined threshold set at 20 per cent of the primary score prior to running the classification. This dual classification allowed for the identification of hybrid interventions while preserving analytical clarity regarding dominant mechanisms.

The classification was implemented using R. Standard preprocessing steps were applied to reduce what Grimmer and Stewart (2013) describe as the "awe-inspiring diversity of language" into a manageable set of features: text was lowercased, punctuation removed, and morphological variation addressed through explicit inclusion of inflected forms in the keyword dictionaries. Algorithmic stemming was not applied, as Swedish morphological complexity and the semantic precision required in a domain-specific corpus made explicit enumeration of inflected forms the more reliable approach. Stop word removal was similarly omitted, as the dictionary-based approach counts only targeted domain-specific terms, and the presence of common function words in the text does not affect scores. Rate normalisation was then applied to ensure results were not biased by document length. As Grimmer and Stewart acknowledge, all quantitative language models are imperfect given the context-dependence of meaning, but what justifies them is their ability to perform a useful social scientific task at a scale that manual analysis cannot (Grimmer & Stewart, 2013).

5.2.3 Strengths and Limitations of Dictionary-Based Classification

Dictionary-based classification offers transparency and theoretical consistency. Because categories are explicitly derived from theory, the link between conceptual framework and empirical classification remains visible and replicable (Grimmer & Stewart, 2013, pp. 284–292). However, the approach is sensitive to wording

variation and may underperform when project descriptions are vague or strategically framed. It also risks misclassification in cases where identical terminology is used across different intervention logics (Grimmer & Stewart, 2013, pp. 280–283). For this reason, the instrument was complemented by external validation steps, as described in Section 5.4.

5.3 Iterative Algorithm Refinement

Following the initial construction of the classification algorithm, an iterative refinement phase was undertaken to improve conceptual alignment and empirical stability. Rather than being fully specified at the outset, the algorithm evolved through successive adjustments prompted by observed inconsistencies in category distributions and misalignment between theoretical expectations and classification output.

5.3.1 Initial Observations and Keyword Refinement

The first round of classification revealed uneven distribution patterns and cases where projects appeared misclassified due to overly broad keyword signals. Generic terms related to “competence”, “training”, and “education” inflated Human Capital scores even when formal qualification or structured education were not central to the intervention.

In response, the keyword dictionaries were revised. This process included (i) adding further context-specific Swedish terms, (ii) removing or excluding overly generic phrases, and (iii) differentiating between strong and weak indicators of intervention logic. For example, generic terms such as “kompetens” (competency) were excluded, while more specific expressions like “yrkesutbildning” (vocational education) were retained and weighted as strong indicators. This stage aimed to reduce superficial matches and improve correspondence between textual signals and underlying causal mechanisms.

5.3.2 Differential Weighting of Keyword Indicators

Following dictionary refinement, a weighted scoring system was introduced for Human Capital, the only category to receive differential weighting for reasons developed in section 5.3.3, as a methodological adaptation developed in response to the observed over-inflation of Human Capital scores from generic terminology. Keywords were divided into strong, medium, and weak indicators, assigned weights of 3, 2, and 1, respectively. This adjustment addressed the problem that not all keywords carry equal conceptual significance. Explicit references to certification or validation were treated as stronger signals than generic references to competence development. The weighted score replaces a simple frequency count and forms the basis for subsequent rate normalisation.

5.3.3 Normalisation and Gating

To control for variation in document length, category scores were normalised as rates per 1,000 words. Projects with fewer than 500 words or 1,000 characters were excluded to ensure a sufficient textual signal. Despite weighting and normalisation, Human Capital continued to dominate in borderline cases due to residual generic terminology. A gating mechanism was therefore introduced. Human Capital classification required either (i) at least three strong indicators combined with a minimum rate threshold, or (ii) a high overall Human Capital rate exceeding a predefined threshold. Projects not meeting these criteria were prevented from being classified as Human Capital. This step substantially reduced over-classification.

The asymmetry between the HC gating mechanism and the treatment of other categories warrants explicit justification. Given that the vocabulary associated with Social Capital and Motivational Activation tends to be expressed implicitly in programme documentation rather than in discrete mechanism-specific terms, one might reasonably ask whether these categories should be assigned compensating lower detection thresholds. The decision not to do so rests on two grounds. First, the vocabulary associated with motivation, coaching, and social connection is

sufficiently generic in Swedish policy documentation that lowering thresholds would substantially increase the risk of false positives, conflating projects that genuinely operate through these mechanisms with projects that use the associated language incidentally. Second, the appropriate response to the implicit expression of these logics is not to lower the evidentiary bar but to be explicit about what the instrument can and cannot recover. Where Social Capital and Motivational Activation signals do not meet the detection threshold, this should be understood as a property of the instrument's capacity rather than evidence that the intervention logic is absent from the project.

5.3.4 Dominance and Mixed Classification

To avoid forced single-category assignment, a dominance ratio between the top two category scores was implemented. Where the ratio fell below a defined threshold of 1.3, the project was classified as Mixed/Hybrid, indicating insufficient dominance of any single logic. A minimum signal threshold of 0.15 per thousand words was also applied. Projects falling below this floor were classified as Other/Unclear. This category reflects a genuine absence of detectable mechanism-specific signal rather than ambiguity between categories. These results replaced manual interpretation with consistent, replicable classification steps.

5.3.5 Calibration of Secondary Analytical Dimension

The functional orientation dimension was developed alongside the primary classification instrument but is calibrated through a different process, and the asymmetry reflects the different analytical roles and theoretical properties of the two dimensions. The primary intervention logic typology was subjected to LLM-assisted external validation because its categories correspond to established theoretical constructs with sufficient cross-contextual grounding to make structured independent classification a meaningful test. The functional orientation categories are more contextually embedded. Derived from the sequential trajectory model

implied by a single national evaluation instrument, their application requires judgments about positional role within a specifically Swedish programme architecture that are less amenable to the same validation design. The dimension was instead calibrated through iterative threshold-setting against repeated manual reading of borderline cases, with all thresholds fixed prior to full corpus application and not adjusted thereafter. This ensures that the classification outputs were not adjusted in response to observed results, which is the core requirement that external validation is designed to guarantee. The functional dimension should accordingly be understood as internally calibrated and descriptively useful rather than externally validated, and its outputs are interpreted in that spirit throughout.

5.4 LLM-Assisted Validation

To assess the external validity of the dictionary-based classification pipeline, a subset of classified project reports was subjected to independent re-classification using a large language model (LLM). The purpose of this validation step was not to refine or alter the algorithm but to test its robustness, specifically, whether the typology's category boundaries were being applied consistently and whether the dominant error patterns were systematic or idiosyncratic.

Classifications were produced using Claude Sonnet 4.6 (Anthropic), an instruction-following language model operating in adaptive mode. Each project report was submitted individually. The model received the full typology as a structured reference table, comprising intervention logic categories, their inclusion criteria, mechanisms, as well as the operationalisation criteria for the Human Capital subtypes. Specifically, the model was asked to assign a primary intervention logic, a secondary intervention logic representing the second most present mechanism, and, when relevant, an appropriate HC subtype. It was also asked to produce a three-sentence written argument justifying each classification. Human Capital was specified as a primary-only category, ineligible as a secondary classification, mirroring the typology gating rule found in the instrument.

Rather than a numerical confidence score, the quality of the LLM's classification was assessed through its written argument. Each three-sentence argument was required to identify the dominant operative mechanisms, name the secondary logic and explain its subordinate status, and locate the classification decision within the typology's inclusion criteria. Arguments that failed to engage with the inclusion criteria, or that mischaracterised the primary mechanisms, were treated as signals of genuine classification uncertainty and prioritised in manual adjudication. This approach treats the LLM's reasoning process as a transparency mechanism rather than a scalar confidence measure, consistent with the qualitative texture of the underlying typology. The LLM validation was conducted after the dictionary pipeline had been finalised, including all threshold parameters, dominance ratio cutoffs, and HC gate criteria. The LLM's classifications were therefore produced against a fixed instrument, and no aspect of the algorithm was modified in response to the validation results.

Primary disagreement was defined as any case in which the LLM's primary classification differed from the algorithm's primary classification. Disagreements were categorised as rank inversions (agreements on logics, disagreement on ordering) or true disagreements (categories differing entirely). Secondary disagreement was defined analogously, restricted to cases where both classifiers assigned a secondary type.

Across 30 sampled projects, drawn randomly from the classified corpus, the classification instrument and LLM agreed on the primary classification in 24 cases, yielding an initial primary agreement rate of 80.0 per cent. The six disagreements comprised three rank inversions and three true disagreements. Rank inversions arose at the Work-First/Human Capital and Work-First/Capability and Health boundaries, both conceptually proximate with direct placement mechanisms, making dominance assessment inherently difficult. All six disagreements were resolved through manual adjudication by the author, in which the original project report was re-read, and a classification decision was defined independently of both

automated outputs. The manual review sided with the algorithm in four of the six cases and with the LLM in two. This yields an effective accuracy of 93.3 per cent (28/30) for the algorithm and 86.7 per cent (26/30) for the LLM against the manual standard.

Secondary disagreement stood at 33.3 per cent (10/30). The dominant pattern among the ten secondary disagreements was the algorithm assigning Work-First as the secondary logic, where the LLM assigned Motivational Activation, a pattern observed in five of the ten discrepant cases. This indicates that the dictionary pipeline captures workplace-placement lexis more readily than coaching or motivational language as a secondary signal, while the LLM reads secondary mechanisms from contextual framing rather than keyword frequency. This asymmetry does not affect primary classifications but represents a relevant limitation for downstream analyses relying on secondary type assignment.

The validation results support the algorithm's primary classifications as substantially reliable. A 93.3 per cent effective accuracy against manual adjudication, combined with the finding that errors are concentrated at conceptually proximate boundaries rather than being randomly distributed, suggests that the classification algorithm is operationally stable and that misclassification results from genuine category ambiguity rather than instrument failure. The validation is not intended as a definitive measure of accuracy; with $n=30$, the figure is better understood as a structured diagnostic of where and how the instrument errs rather than as a precise population-level estimate.

6. Results

Chapter 6 presents the classification results across the full corpus of 484 European Social Fund Programme Area 2 final project reports. Section 6.1 reports the overall distribution of primary intervention logics and addresses the residual categories. Section 6.2 presents the subtype analysis of the 103 projects classified as primary

Human Capital, examining internal variation in the form of competence development represented across the portfolio. Section 6.3 reports the distribution of secondary intervention logics and examines the structural patterns through which mechanisms are combined within individual projects. Section 6.4 presents the functional orientation results, capturing each project's positional role within the labour market trajectory independently of its primary logic classification. Section 6.5 reports the dominance ratio and signal strength analysis, providing a measure of classificatory certainty across categories. Together, these five dimensions constitute the empirical foundation on which the discussion in Chapter 7 is built.

6.1 Overall Distribution of Primary Intervention Logics

The classification procedure was applied to the full corpus of 484 European Social Fund Programme Area 2 final project reports. Each project was assigned a single primary intervention logic based on the highest-scoring typology category, subject to the gating and dominance rules described in Chapter 5.

6.1.1 Primary Classification Results

The results reveal a portfolio that is neither evenly distributed across intervention logics nor reducible to a single dominant approach. Seven outcome categories are represented in the final classification, comprising five substantive intervention logics and two residual categories.

Table 2: Distribution of Primary Intervention Logics (n=484)

<i>Primary intervention logic</i>	<i>n</i>	<i>%</i>
<i>Work-First</i>	131	27.1
<i>Human Capital</i>	103	21.3
<i>Mixed/Hybrid</i>	76	15.7
<i>Capability/Health</i>	72	14.9
<i>Motivational Activation</i>	62	12.8
<i>Other/Unclear</i>	21	4.3

Work-First and Human Capital are the two most frequently occurring primary categories, together accounting for 48.4 per cent of all projects in the sample. This concentration indicates that a majority of Programme Area 2-funded interventions are oriented toward either rapid labour market attachment through workplace exposure or structured investment in participants' productive capacities, the two paradigms most consistently represented in the broader active labour market policy literature (Card et al., 2018; Kluve, 2010). The remaining substantive categories (Capability and Health, Motivational Activation, and Social Capital) account collectively for 153 projects (31.6 per cent) of the portfolio, indicating that while alternative intervention logics are present, they occupy more varied positions relative to the two dominant paradigms.

6.1.2 Residual Categories and the Prevalence of Hybrid Logics

Two features of the distribution warrant particular attention. The first is the size of the Mixed/Hybrid category. At 76 projects, or 15.7 per cent of the sample, more than one in seven classified projects does not exhibit a single clearly dominant intervention logic under the applied dominance threshold. This is not a methodological artefact but a substantively meaningful finding. It indicates that a non-negligible share of Programme Area 2 projects is designed around combinations of mechanisms rather than a single coherent theory of change. As discussed in Chapter 3, hybrid intervention forms are an expected feature of labour market policy design rather than anomalies, and their prevalence here is consistent with the view that integration is often treated as a multidimensional process requiring coordinated responses to overlapping barriers.

The second noteworthy feature is the gap in absolute terms between the largest and smallest substantive categories. Work-First accounts for 131 projects while Social

Capital accounts for 19, a difference of 112 projects. The Mixed/Hybrid category alone outnumbers Social Capital and is comparable in size to Capability and Health (72). This disparity suggests that certain problem representations and theories of change are structurally more likely to receive funding under Programme Area 2 than others, a point returned to in the discussion in Chapter 7.

It is worth acknowledging transparently that the Other/Unclear (n=21, 4.3 per cent) category is, in part, a methodological artefact. It captures not only projects that genuinely lack a dominant intervention logic but also projects whose logic is expressed in vocabulary that the dictionary does not recover. The two cannot be fully distinguished from the outside. What can be said is that the 4.3 per cent rate is low enough that it is unlikely to materially distort the distributional findings, and that retaining the category as an explicit residual is preferable to absorbing these cases into substantive classifications where they do not belong. Following Bowker and Star (1999), the residual categories are not failures of the instrument but features of it, making visible what the typology's boundaries cannot cleanly contain, whether through insufficient signal density or through distributions that resist the dominance logic the instrument applies.

6.2 Human Capital Subtypes

The 103 projects classified as primary Human Capital were subject to a second round of classification to identify the dominant form of competence development represented in each project description. This subtype analysis is motivated by the theoretical heterogeneity within the Human Capital category itself. As established in Chapter 3, human capital interventions may target formal qualifications, language acquisition, the recognition of existing competencies, or structured vocational training. These subtypes rest on partially distinct assumptions about what kind of skill deficit is being addressed and through what mechanism it can be remedied.

6.2.1 Human Capital Subtype Results

The distribution of Human Capital subtypes reveals a pronounced internal concentration.

Table 3: Distribution of Human Capital Subtypes (n=103)

<i>Human Capital subtype</i>	<i>n</i>	<i>%</i>
<i>HC1 Credential</i>	3	2.9
<i>HC2 Language</i>	50	48.5
<i>HC3 Validation</i>	14	13.9
<i>HC4 Formal/Vocational Education</i>	36	35.0

Language-focused interventions (HC2) account for 50 projects, or 48.5 per cent of all Human Capital cases, making them the most common subtype and the single largest coherent sub-category in the entire dataset. Formal or Vocational projects (HC4) constitute a further 35.0 per cent, meaning that these two subtypes together account for over 83 per cent of all Human Capital classifications. The remaining two present subtypes, Credential (HC1) and Validation of Prior Learning (HC3), collectively account for 17 projects, or 16.5 per cent of the Human Capital category. HC5 (Other Structured Training, the residual subtype assigned when no subtype-specific keywords are detected) returned no cases in this corpus, indicating that all Human Capital projects contained sufficient subtype-specific vocabulary to be assigned to one of the four substantive categories.

6.2.2 The Dominance of Language Interventions

The concentration of Human Capital projects with the language subtype is statistically significant and warrants interpretation beyond mere distributional observations. It is unlikely to be coincidental. Programme Area 2's eligibility criteria explicitly included newly arrived migrants as a target group, and language acquisition (most commonly Swedish for Immigrants/SFI or occupationally oriented language training) is recognised as one of the most clearly defined institutional pathways for this population. The prominence of HC2 thus reflects, at

least in part, the demographic composition of the target population and the institutional framing of linguistic competence as a primary barrier to labour market participation for this group.

This finding also has implications for how the Human Capital category should be interpreted. Together, language training and formal vocational education account for over 80 per cent of Human Capital classifications, indicating that the portfolio's Human Capital logic is oriented primarily toward foundational competence acquisition rather than credential production. Language training programmes may operate through distinct mechanisms and exhibit different effect profiles than general vocational education, and both differ considerably from the credential-oriented investment that the human capital literature most directly associates with strong medium-to-long-term labour market returns.

6.2.3 Formal Vocational Education and Validation Interventions

The second largest subtype, formal and vocational education (HC4, n=36), represents projects organised around structured curriculum-based training without a primary emphasis on formal certification. These projects include adult education through Komvux (municipal adult education), vocational programmes, apprenticeships, and preparatory vocational courses. At 35.0 per cent of Human Capital classifications, HC4 constitutes a substantial and substantively distinct component of the portfolio. The mechanism assumed to generate change in these projects is structured skill acquisition through defined educational pathways, which more closely approximate the canonical human capital investment model described in section 3.3.1 than either language training or credential-focused programmes.

Validation of Prior Learning (HC3, n=14) represents the third subtype at 13.9 per cent, encompassing projects whose primary activity involved the formal recognition of competencies acquired through work experience or non-formal education. Despite its growing policy relevance in the Swedish context, validation

occupies a modest share of the Human Capital portfolio. The low frequency may reflect the relative difficulty of implementing validation processes within the time and resource constraints of project-based funding, or the comparatively recent institutionalisation of validation as a policy instrument. Its presence nonetheless points to a mechanism that addresses competence recognition rather than competence formation, and which is of relevance for participants with internationally acquired qualifications that are not formally recognised in the Swedish labour market.

Credential-oriented interventions (HC1, $n=3$) account for only 2.9 per cent of Human Capital classifications, indicating that projects explicitly oriented toward the acquisition of formal qualifications, certification, or professional licenses are rare within the portfolio. This is a striking finding given the prominence of certification in labour market policy discourse, and it suggests that the Programme Area 2 portfolio, despite its substantial Human Capital share, is not primarily oriented toward producing formal credentials.

6.2.4 Methodological and Substantive Implications of Subtype Concentration

The subtype distribution carries both methodological and substantive implications. Methodologically, it demonstrates that the frequency-based classification procedure successfully differentiates between qualitatively distinct forms of human capital investment, rather than collapsing all training-related language into a single undifferentiated category. The substantial variation across subtypes, from 48.5 per cent for language training to 2.9 per cent for credential-oriented projects, confirms that the subtype instrument is detecting meaningful differences in how projects describe their core competence development activity rather than simply distributing projects arbitrarily.

Substantively, it indicates that what the Programme Area 2 portfolio funds under a human capital logic is not a uniform type of skill investment but a set of interventions varying considerably in their target competency, their institutional context, and by extension their likely effect profiles. The dominant combination of language training and formal vocational education reflects a portfolio oriented primarily toward foundational competence acquisition for populations far from the labour market, particularly newly arrived immigrants and individuals with limited formal qualifications. Credential acquisition, which the literature most directly associates with improved medium-to-long-term labour market outcomes, is effectively absent.

6.3 Secondary Intervention Logics

The classification procedure allowed for the assignment of a secondary intervention logic where a project exhibited a clearly defined primary category and sufficient signal strength in a second dimension. Secondary classification was restricted to projects assigned a substantive primary logic, excluding Mixed/Hybrid and Other/Unclear cases on the grounds that these categories already reflect distributional ambiguity rather than analytical clarity. This yields 387 projects eligible for secondary classification.

6.3.1 Secondary Classification Results

The distribution of secondary logics differs markedly from the primary distribution and reveals a distinct structural pattern in how intervention mechanisms are combined within individual projects.

Table 4: Distribution of Secondary Intervention Logics (n=387)

<i>Secondary logic</i>	<i>n</i>	<i>%</i>
<i>Work-First</i>	96	24.8
<i>Motivational Activation</i>	72	18.6
<i>Capability/Health</i>	62	16.0

<i>Social Capital</i>	21	5.4
<i>Human Capital</i>	0	0
<i>None</i>	136	35.1

Note: Percentages are rounded to one decimal place and may not sum to 100

Work-First emerges as the most common secondary logic, appearing in 96 projects or 24.8 per cent of those eligible for secondary classification. Motivational Activation (n=72, 18.6 per cent) and Capability and Health (n=62, 16.0 per cent) constitute the second and third most frequent secondary logics, respectively, while Social Capital appears as a secondary dimension in 21 projects (5.4 per cent). A substantial majority of eligible projects (n=136, 35.1 per cent) received no secondary classification, indicating that their textual content was sufficiently concentrated in a single dimension to preclude meaningful secondary signal detection.

6.3.2 Work-First as a Structuring Secondary Logic

The prominence of Work-First as a secondary logic is the most striking feature of this distribution, and it warrants careful interpretation. It suggests that workplace exposure, through internship, wage subsidy, or direct job entry, functions not only as a standalone intervention strategy but as a recurrent complementary component within projects primarily oriented toward other logics. In a number of cases, projects whose primary mechanisms are skill formation, health rehabilitation, or motivational support also incorporate a work-first element, whether as a culminating phase of the intervention, as a parallel track, or as a contextual setting within which other mechanisms operate.

This pattern is consistent with the sequential understanding of labour market integration developed in Chapter 3, where different mechanisms may operate at different stages of a participant's trajectory. A project classified as primarily Human Capital, for instance, may embed workplace placement as a secondary mechanism through which newly acquired skills are consolidated or demonstrated to potential

employers. Similarly, a Capability and Health project may incorporate supported employment as a component of the rehabilitation process rather than its primary logic. The relative prominence of Work-First as a secondary dimension, therefore, does not simply indicate that projects are poorly focused; it may instead reflect a deliberate design logic in which labour market attachment is treated as both an intermediate step and an outcome.

6.3.3 Secondary Logics and Portfolio Hybridity

Taken together, the secondary classification results indicate that the majority of Programme Area 2 projects with a clearly defined primary logic are textually concentrated to a single intervention dimension, as 35.1 per cent of projects received no secondary classification under the applied threshold. This reflects the conservative nature of the secondary assignment rule, which requires the second-ranked category to score at least 20 per cent of the top-scoring category. Among those projects where a secondary logic is detectable, Work-First is the most common complementary dimension, followed by Capability and Health, as well as Motivational Activation. This pattern suggests that where project designers do combine mechanisms, they tend to sequence or embed labour market attachment, health-related support, or motivational elements alongside a dominant primary logic. The secondary dimension thus provides an important corrective to any impression of categorical neatness, while the high rate of single-dimension concentration confirms that primary classification captures the substantive core of most projects in the corpus. This distinction between primary classification and complement mechanism homogeneity matters for how the results are read and is addressed further in the discussion of the instrument's limitations in section 8.4.

6.4 Functional Orientation Within the Labour Market Trajectory

In addition to the intervention logic classification, each project was assigned a functional orientation reflecting its positional role within the labour market trajectory. This dimension is analytically independent of the typological classification; two projects sharing the same primary intervention logic may differ substantially in whether they are oriented toward stabilising participants prior to activation, preparing them for labour market entry, or facilitating direct transition into employment. The functional dimension is, therefore, able to capture a complementary aspect of intervention design that the logic-based typology does not fully address.

6.4.1 Distribution of Functional Orientations

The distribution of functional orientations reveals that the majority of projects operate across multiple stages of the labour market trajectory simultaneously.

Table 5: Distribution of Functional Orientations (n=484)

<i>Function</i>	<i>n</i>	<i>%</i>
<i>Mixed/Multiple</i>	226	46.7
<i>Direct Labour Market Entry</i>	154	31.8
<i>Other</i>	74	15.3
<i>Stabilising</i>	26	5.4
<i>Assessment</i>	3	0.6
<i>Preparatory</i>	1	0.2

Mixed or multiple functional orientations constitute the largest single category, accounting for 226 projects or 46.7 per cent of the sample. Direct Labour Market Entry is the second largest category at 154 projects (31.8 per cent), indicating that over a third of projects are primarily oriented toward facilitating immediate transitions into employment. Stabilising functions, encompassing interventions aimed at establishing basic routines, social participation, or psychological readiness

prior to more direct activation, account for 26 projects (5.4 per cent), while Assessment-oriented projects, whose primary purpose is the evaluation of participants' work capacity or competencies, represent only 3 projects (0.6 per cent). The Preparatory category returned only one project as its sole designated function, which does not necessarily indicate that preparatory functions are nearly absent from the portfolio. Where preparatory signals co-occur with other functional flags, projects are assigned Mixed/Multiple rather than Preparatory independently, meaning preparatory orientation is likely present as a component of more complex functional profiles. The remaining 74 projects (15.3 per cent) were assigned to the Other category, capturing functional orientations that do not map cleanly onto the predefined dimensions.

6.4.2 Functional Orientation by Primary Logic

The two largest categories (Mixed/Multiple and Direct Labour Market Entry) together account for 78.5 per cent of the sample. The remaining four categories collectively account for 21.5 per cent.

Table 6: Functional Orientation by Primary Intervention Logic

<i>Primary type</i>	<i>Mixed/Multiple</i>	<i>Direct Entry</i>	<i>Other</i>	<i>Stabilising</i>	<i>Assessment</i>	<i>Preparatory</i>
<i>Work-First</i>	42.7	55.0	2.3	0	0	0
<i>Human Capital</i>	49.5	38.8	8.7	2.9	0	0
<i>Capability/Health</i>	61.1	9.7	9.7	19.4	1.4	0
<i>Mixed/Hybrid</i>	61.8	25.0	9.2	3.9	0	0
<i>Motivational Activation</i>	35.5	19.4	32.3	8.1	4.8	0
<i>Social Capital</i>	31.6	21.2	36.8	5.3	0	5.3

Note: Values are row percentages. Rows may not sum to 100 due to rounding.

The cross-tabulation reveals that functional orientation varies systematically across primary logics rather than being randomly distributed. Work-First projects are heavily concentrated in Direct Labour Market Entry (55.0 per cent) and are completely absent from all other functional orientations as the primary category,

consistent with their emphasis on rapid attachment. Capability and Health projects are the only category where Stabilising constitutes a substantive share (19.4 per cent), reflecting the health-related barriers that must be addressed before direct activation becomes feasible, and consistent with the capability approach's location of disadvantage in reduced capacity to convert opportunities into participation rather than in skill deficits or placement barriers alone. The relatively high Mixed/Multiple share further suggests that health-related interventions frequently combine stabilising functions with preparatory or direct entry components, consistent with the sequential design logic discussed in Chapter 3. Social Capital projects show the most dispersed profile, with over a third falling into the Other category, suggesting their functional orientation resists the sequential trajectory model. These patterns are interpreted in relation to the research questions in section 7.3.

6.5 Classification Certainty and Signal Strength

Alongside the distributional results, the dominance ratio provides a measure of how clearly defined each primary classification is. It is calculated as the ratio of the highest-scoring category to the second-highest. A high dominance ratio indicates that one intervention logic substantially outperforms all others in each given project, while a low ratio indicates proximity between the top two categories and consequently greater classificatory uncertainty.

6.5.1 Dominance Ratios by Primary Category

Table 7: Median Dominance Ratio by Primary Category

<i>Primary Category</i>	<i>Median Dominance</i>
<i>Human Capital</i>	7.32
<i>Work-First</i>	2.82
<i>Social Capital</i>	2.62
<i>Capability/Health</i>	2.28
<i>Motivational Activation</i>	2.00

Human Capital exhibits by far the highest median dominance ratio at 7.32, indicating that when a project is classified as Human Capital, its top-scoring category substantially outscores the second. This is a direct consequence of the weighted scoring and gating mechanisms applied to this category. Projects must contain a sufficient density of strong indicators to receive a Human Capital classification, meaning that those that do tend to be textually concentrated around human capital signals to a degree that substantially exceeds any competing dimension.

Work-First exhibits an intermediate dominance ratio of 2.82, while Social Capital, Capability and Health, and Motivational Activation cluster between 2.00 and 2.62. Mixed/Hybrid projects sit just below the dominance threshold of 1.3, with a median of 1.12. This is consistent with the category being assigned precisely when the ratio between the top two categories falls below that threshold.

Table 8: Mean Top-Scoring Category Rate by Primary Intervention Logic

<i>Primary intervention</i>	<i>Mean top1</i>
<i>Human Capital</i>	20.39
<i>Work-First</i>	3.45
<i>Capability/Health</i>	3.25
<i>Social Capital</i>	2.92
<i>Motivational Activation</i>	2.18
<i>Mixed/Hybrid</i>	2.00
<i>Other/Unclear</i>	0.00

The variation in mean top-scoring values across categories is equally informative. Human Capital projects exhibit a mean top score of approximately 20.39, substantially higher than the rest. Other/Unclear projects have a mean top score of zero, confirming that their residual classification reflects a genuine absence of detectable signal rather than a marginal failure to meet category thresholds.

6.5.2 Methodological Implications of Dominance Variation

The substantial variation in dominance ratios and signal strength across categories reflects genuine differences in how intervention logics are expressed in project documentation and has methodological implications for the interpretation of classification results. Human Capital classifications carry a higher degree of textual certainty than classifications in other categories, given the stringency of the gating mechanism and the corresponding concentration of signal required to trigger it. Classifications in categories such as Motivational Activation or Social Capital, where dominance ratios are lower and mean scores more modest, should be interpreted as indicating the relative prominence of a given logic rather than its exclusive presence.

This variation does not undermine the validity of the classification procedure, but it does underscore the importance of treating the typology as a structured analytical tool rather than as a set of discrete and mutually exclusive categories. The dominance and signal strength data reported here provide the empirical basis for calibrating the interpretive confidence attached to different parts of the classification output and inform the discussion of methodological limitations in Chapter 7.

7. Discussion

The discussion that follows interprets the results presented in Chapter 6 in direct relation to the research questions posed in Chapter 1, organised in the order in which those questions were raised. Section 7.1 addresses RQ1 by assembling a composite characterisation of the portfolio across five analytical dimensions, establishing what the European Social Fund Programme Area 2 project corpus reveals about the intervention logic composition of Swedish labour market activation. Section 7.2 deepens that characterisation by examining the relationship between the language of project documentation and the mechanisms identified through classification,

speaking to both RQ1 and the sub-RQ by showing how administrative categories obscure meaningful programmatic variation. Section 7.3 addresses the sub-RQ directly, interpreting the concentration, hybridity, and functional orientation findings in relation to the governance conditions that produce them. Section 7.4 addresses RQ2, reflecting on the extent to which the classification instrument reliably distinguishes between distinct intervention logics and what the dominance and signal-strength analyses reveal about the boundaries of its discriminatory capacity.

7.1 The Intervention Logic Composition of Programme Area 2

The European Social Fund Programme Area 2 portfolio cannot be adequately characterised through its primary intervention logic distribution alone. This section addresses RQ1: *How can Swedish Labour market activation projects be characterised in terms of their underlying intervention logics?* A composite reading across five analytical dimensions: primary logic distribution, Human Capital subtype, secondary logic combinations, functional orientation, and dominance patterns, reveals a portfolio that is more internally differentiated and more analytically interesting than aggregate descriptions imply.

The primary distribution establishes the overall composition of the portfolio. Work-First and Human Capital are the predominant logics, together accounting for just under half of all classified projects. Mixed/Hybrid constitutes a significant third, showing that more than one in seven projects does not revolve around a single coherent theory of change. Capability and Health, and Motivational Activation, occupy intermediate positions, while Social Capital is the smallest substantive category. The portfolio is not dominated by a single approach nor evenly spread across the typology; instead, it is built around two main paradigms with a secondary layer of more diverse approaches.

The subtype analysis of Human Capital projects adds the first layer of internal differentiation. Language-oriented projects and formal or vocational education together account for over 83 per cent of Human Capital classifications, with credential-oriented projects (those most directly approximating the canonical human capital investment model) accounting for only 2.9 per cent. What the portfolio funds under a Human Capital logic are, therefore, not a uniform type of skill investment but a set of interventions oriented primarily toward foundational competence acquisition for populations far from the labour market. This internal differentiation is not merely a classification refinement; it has direct implications for how the portfolio's effectiveness should be assessed, since language training, vocational education, and credential production each carry different theoretical expectations about the timing and magnitude of labour market returns.

The secondary classification adds a further dimension by revealing how intervention logics are combined within individual projects. Work-First is the most common secondary logic, followed by Motivational Activation and Capability and Health, while just over a third of eligible projects received no secondary classification at all, indicating that their primary logic captures the dominant and near-exclusive mechanism expressed in the documentation. Where secondary signals are present, the pattern is structured rather than random; Work-First appears as a complementary mechanism across projects primarily organised around skill formation, health rehabilitation, and motivational support, suggesting a sequential design logic in which labour market attachment functions as both an intermediate step and an outcome. Human Capital receives no secondary classifications across the eligible portfolio, consistent with its high signal density and the stringency of the gating mechanism. This asymmetry between Human Capital as a primary logic that dominates its projects completely, and Work-First as a secondary logic that pervades projects primarily organised around other mechanisms, suggests that the two dominant paradigms operate differently within the portfolio: one as a self-contained and textually exclusive logic, the other as a structuring component that recurs across diverse intervention types.

A fourth layer is added by the functional orientation dimension, which captures where each logic positions participants within the labour market trajectory, independently of the causal mechanisms it deploys. Work-First projects are heavily concentrated in Direct Labour Market Entry, consistent with their emphasis on rapid attachment. Human Capital projects are distributed across Mixed/Multiple and Direct Entry, reflecting a logic oriented toward building capacity that feeds into employment transitions. Capability and Health is the only primary logic where Stabilising constitutes a substantive share, reflecting the health-related barriers that must be addressed before more direct activation becomes feasible. Motivational Activation projects distribute notably across Other and Mixed/Multiple, indicating that their functional orientation is diffuse in ways that resist clean positional assignment. Social Capital projects show the most dispersed functional profile, with over a third falling into the Other category, suggesting that the relational mechanisms through which this logic operates are genuinely orthogonal to the sequential trajectory model; they do not stabilise, prepare, or place participants in any straightforward sense, but restructure the network conditions within which labour market access becomes possible.

A fifth and final layer is provided by the dominance ratio analysis, which captures not only what logics are present but how coherently and exclusively each logic is expressed within the projects it dominates. Human Capital exhibits by far the highest median dominance ratio, meaning that when a project is classified as Human Capital, its documentation is organised around that logic to a degree that substantially outscores all competing dimensions. Work-First sits at an intermediate level, dominant but with more competing signals present. Capability and Health, Social Capital, and Motivational Activation occupy a lower intermediate range, indicating that these logics are detectable and prominent within their classified projects but rarely exclusive. This variation means that the portfolio's logics are not held with equal coherence across the typology; some projects are organised around

a single mechanism with high textual concentration, others around a dominant logic that coexists with meaningful competing signals.

Taken together, these five layers characterise the European Social Fund's Programme Area 2 portfolio as a structured but internally differentiated activation landscape. It is dominated at the primary level by Work-First and Human Capital, but that dominance conceals substantial internal variation in what those logics involve in practice, how they combine with secondary mechanisms, where they position participants in the labour market trajectory, and how coherently they are expressed in project documentation. The characterisation that emerges is therefore not a portfolio organised around two paradigms. Rather, it is a portfolio in which two paradigms provide the structural frame within which a considerably more varied set of intervention configurations operate. These configurations differ not only in their dominant logic but in how that logic is held, combined, and positioned within the activation trajectory.

7.2 The Discrepancy Between Rhetoric and Classified Mechanisms

This section addresses both RQ1 and the sub-RQ, though from different angles. It deepens the characterisation of the portfolio established in 7.1 by examining the relationship between the language of project documentation and the intervention logics identified through classification, speaking to *RQ1: How can Swedish labour market activation projects be characterised in terms of their underlying intervention logics?* and to the *Sub-RQ: What does this characterisation reveal about the governance context that shapes activation policy?* by showing how the governance architecture shapes what gets expressed and counted in that documentation.

The iterative development of the classification instrument was substantially driven by a persistent discrepancy. Generic references to competence development, skills,

and training were pervasive across the corpus, yet their presence in a document was not, on its own, a reliable indicator of Human Capital logic in any analytically meaningful sense. The introduction of weighted scoring and gating mechanisms, requiring strong indicators and not just training-related vocabulary, substantially reduced the Human Capital classification rate from what an initial pass would have yielded. This discrepancy is itself analytically informative, as it illustrates precisely why characterising activation projects through administrative labels alone produces a misleading account of the mechanism through which change is assumed to occur.

The same activity, in this case training, can serve fundamentally different purposes depending on the underlying theory of change. It may function as genuine skill investment in the Human Capital sense, as a mechanism for building daily routines and psychological readiness in the Motivational Activation sense, as a vehicle for expanding social contact with employers and professional communities in the Social Capital sense, or as an ancillary component of a Work-First trajectory in which the primary mechanism is workplace exposure rather than skill accumulation. The analytical move required to distinguish between these is to look past the activity label to the embedded causal assumption, which is precisely what the gating mechanism was designed to do.

What the results reveal, when read through this framework, is that educational and training language functions partly as a form of symbolic legitimacy in project documentation. References to skills development carry strong positive normative associations, with long-term investment, individual empowerment, and alignment with the knowledge-based economy discourse that shapes European Social Fund programmes. Drawing on Bowker and Star's (1999) account of the politics of visibility, what gets named and counted is not neutral but reflects what the institutional environment treats as legitimate and reportable. Project owners who frame their interventions in training terms are, in part, positioning themselves within a discursive field that the funding system rewards, regardless of whether the underlying mechanism of change is primarily based on skill accumulation. The

vocabulary of human capital investment has, in this sense, become decoupled from its theoretical referent in a substantial portion of the corpus. The broader governance implications of this dynamic, how the infrastructure of the European Social Fund's infrastructure shapes not only what gets funded but what can be expressed and made legible, are developed in section 7.3.

The subtype analysis sharpens this point considerably. Of all Human Capital classifications, language-oriented and formal vocational education projects together account for over 83 per cent, while credential-oriented projects, those most directly approximating the canonical human capital investment model, account for only 2.9 per cent. The aggregate Human Capital share of the portfolio, therefore, substantially overstates the degree to which Programme Area 2 funds the kind of credential-producing skill investment that the human capital tradition most directly associates with strong medium-to-long-term labour market returns. What dominates instead is language acquisition and structured vocational training. These are forms of competence development whose institutional contexts, target populations, and expected effect profile differ considerably from one another and from the credential model, yet which appear under the same administrative and typological label.

This finding has direct practical implications for how the portfolio is evaluated. Assessments of the European Social Fund's Programme Area 2 that treat the Human Capital category as analytically homogenous will aggregate language training, credential programmes, validation of prior learning, and vocational education into a single effectiveness estimate, obscuring variation that is theoretically meaningful and empirically consequential. The subtype classification developed in this thesis provides the granularity needed to disaggregate these interventions and assess them against the theoretical expectations most appropriate to each logic. More broadly, the discrepancy between rhetorical and classified mechanisms demonstrates that the activity categories through which activation portfolios are typically described are insufficient analytical tools. This is not because they are inaccurate descriptions of what projects do, but because they are silent on how and why those activities are

expected to produce change. It is precisely this silence that a mechanism-based classification is designed to break.

7.3 Governance, Concentration, and the Architecture of Activation Policy

This section addresses the sub-RQ directly: *What does this characterisation of intervention logics reveal about the governance context that shapes activation policy?* The answer that emerges from the results is threefold. The distribution of intervention logics across the portfolio reflects institutional selection pressures embedded in the European Social Fund governance architecture. The hybridity revealed by secondary classification and dominance patterns indicates that activation policy operates in practice as a configuration of interacting mechanisms rather than a set of discrete programme types. And the functional orientation results reveal that different logics systematically position participants at different points in the labour market trajectory, producing a portfolio that processes different populations through structurally distinct pathways. Together, these three observations constitute an answer to the sub-RQ that goes considerably beyond what the primary distribution alone could support. The governance conditions referred to throughout this discussion encompass the funding criteria, reporting requirements, output indicators, and accountability frameworks through which the European Social Fund structures what kinds of interventions can be proposed, funded, and counted.

The concentration of Work-First and Human Capital logics, and the corresponding marginalisation of Social Capital and the comparatively modest presence of Motivational Activation, are not random outcomes. Before developing an explanation, however, the most obvious alternative account warrants direct acknowledgement. Work-First and Human Capital interventions may dominate the portfolio for reasons that are primarily practical rather than institutional: they are operationally simpler to implement at scale, they draw on well-established delivery

models with which project owners are already familiar, and they generate outputs that are straightforwardly observable within a project cycle. On this account, the concentration pattern would reflect the organisational capacities and implementation preferences of project owners rather than the selection pressures of the funding architecture.

This explanation is not false, but it is insufficient on its own. Practical scalability could account for the dominance of Work-First and Human Capital over Capability and Health; it cannot as readily account for the near-absence of Social Capital, whose relational mechanisms are not inherently more resource-intensive or technically complex than other logics, but whose effects do not translate into the indicator vocabulary the programme uses to count success. That asymmetry is what a governance legibility account explains, and a purely practical account leaves unaddressed. The European Social Fund's funding criteria, standardised reporting templates, and output-focused evaluation frameworks are not neutral administrative instruments. Rather, they structure which kinds of interventions are legible as legitimate policy responses, which activities can be counted and reported, and therefore which kinds of projects are most likely to be proposed and funded (Bowker & Star, 1999). Work-First and Human Capital interventions are, in this sense, the most institutionally legible forms of labour market activation. They map cleanly onto established administrative categories, generate clearly measurable short-term outputs such as employment entries or qualifications obtained, and align with the dominant diagnostic frameworks through which unemployment is framed at the programme level. The governance legibility argument does not require the practical alternative to be false; the two mechanisms are likely to operate simultaneously and to be empirically inseparable from distributive data alone. What the governance framing adds is explanatory traction for the specific shape of the concentration, not simply that some logics are more common than others, but that the logics most present are precisely those whose outputs are most legible within the European Social Fund's monitoring infrastructure, while those most marginal

are precisely those whose effects are relational, attitudinal, or cumulative in ways that resist individual-level output reporting.

The marginal position of Social Capital warrants interpretation that goes beyond its face value as a distributional finding. Social capital formation is relational and diffuse; its effects operate at the network and community level in ways that are not easily reducible to the individual-level output indicators that European Social Fund monitoring requires. The low classification frequency of this logic may therefore reflect not simply a funding preference for legible mechanisms, but a reporting architecture that actively renders certain theories of change invisible, not by excluding them explicitly but by providing no adequate language in which they can be expressed and counted. Motivational Activation occupies a different position. At 12.8 per cent of the portfolio, coaching and behavioural-readiness language is sufficiently expressible within European Social Fund documentation to generate detectable signals. Yet, its presence remains modest relative to its theoretical coherence and empirical grounding, consistent with the view that mechanisms expressed through attitudinal register rather than discrete vocabulary remain harder to operationalise within a target-driven reporting architecture.

The hybridity revealed by the secondary classification and dominance analyses adds a further governance dimension. The finding that Work-First functions not only as the most frequent primary logic but also as the most common secondary dimension across the portfolio, appearing as a complementary mechanism in projects primarily organised around skill formation, health rehabilitation, and motivational support, is consistent with a sequential activation design in which labour market attachment is treated as both an intermediate step and an outcome. This pattern suggests that the governance architecture does not simply privilege certain logics over others; it also shapes how logics are combined. Where funding systems consistently reward direct labour market entry as the primary measurable output, project designers have an incentive to embed Work-First elements as a downstream component even within projects whose primary mechanism operates

through a different causal pathway. The prevalence of Work-First as a secondary logic may therefore be partly a product of the same accountability infrastructure that produces its dominance as a primary logic. A governance system that makes labour market attachment the default measure of success will tend to produce projects that orient toward that outcome regardless of their primary mechanism.

The functional orientation results provide the final layer of the sub-RQ answer. The cross-tabulation between primary logic and functional orientation reveals that different logics do not simply vary in their causal assumptions; they systematically position participants at different points along the labour market trajectory. Work-First projects are concentrated in Direct Labour Market Entry, Capability and Health projects span Stabilising and Mixed/Multiple, and Social Capital projects resist the sequential model altogether. The sequential trajectory model assumes that activation follows a broadly linear logic (stabilisation precedes preparation, preparation precedes entry), and the functional orientation dimension is built around that assumption. That Social Capital projects resist this categorisation reflects something substantive about how relational mechanisms operate. Network formation and bridging tie development do not follow a staged temporal logic in the way that skill acquisition or health rehabilitation do, but are diffuse and cumulative. A governance architecture organised around sequential activation trajectories and individual-level output indicators is therefore structurally ill-equipped to accommodate Social Capital logics, not only because their outputs are hard to count but because their underlying temporal and social logic is incommensurable with the sequential model the architecture assumes. This adds a further dimension to the governance argument developed above: the marginalisation of Social Capital may reflect not only the illegibility of relational outputs but the incompatibility of relational mechanisms with the programme model itself.

Before drawing governance conclusions from these positional patterns, one caveat applies. The functional orientation categories are derived from a specifically

Swedish sequential activation model and calibrated internally against this corpus, meaning the findings describe how projects in this programme position themselves relative to a sequential labour market trajectory, not how they would score against a universal standard. This does not affect their interpretive value for the sub-RQ. The governance argument requires only that the positional patterns are systematically structured rather than randomly distributed, and that condition is met. This structural differentiation has governance implications that extend beyond the individual project level. A portfolio in which logics specifically designed to address functional, attitudinal, and relational barriers (those most prevalent among participants furthest from the labour market) occupy comparatively marginal positions raises questions about its calibration to the needs of its most disadvantaged target groups. These are precisely the participants for whom activation readiness cannot be assumed and for whom Capability and Health, Motivational Activation, and Social Capital interventions are most directly relevant. The functional orientation results do not demonstrate this misalignment directly, but they provide the empirical basis on which that question can be posed systematically rather than incidentally, which is itself a contribution of the typology.

Taken together, the concentration pattern, the hybridity finding, and the functional orientation results produce a characterisation of the European Social Fund's Programme Area 2 that is considerably more analytically tractable than aggregate programme descriptions allow. The portfolio reflects a governance architecture that privileges institutional legibility, rewards measurable attachment outcomes, and creates structural conditions in which certain theories of change are consistently more expressible, fundable, and detectable than others. Understanding activation policy in terms of its intervention logic composition, rather than simply its aggregate outputs, makes these structural conditions visible and opens them to the kind of systematic scrutiny that the present typology was designed to enable.

7.4 The Reliability of the Classification Instrument

This section addresses RQ2 directly: *To what extent can a dictionary-based, theory-driven classification instrument reliably distinguish between distinct intervention logics in a large Swedish-language policy corpus?* The evidence supports a positive but carefully calibrated answer. The instrument performs reliably at the primary classification level, with confidence that varies systematically across categories in ways that the analysis can characterise, and less reliably at the secondary level in ways that are diagnostically informative rather than simply limiting.

The LLM-assisted validation reported in section 5.4 provides the most direct evidence bearing on this question. At $n=30$, the sample is not designed to yield a precise population accuracy estimate, and the 93.3 per cent figure should be read accordingly: as a well-grounded approximation whose primary inferential value lies not in the point estimate itself, but in the pattern of disagreement it reveals. Crucially, the six initial disagreements were not randomly distributed; three were rank inversions at conceptually proximate boundaries, and three were substantive disagreements. An instrument that errs most frequently at theoretically adjacent margins behaves as expected, given a typology of overlapping constructs, and produces the kind of structured error profile that can be characterised and accounted for in interpretation. Whether that error rate holds precisely at 93.3 per cent across the full corpus cannot be confirmed from a sample of this size. What can be confirmed is that misclassification, when it occurs, reflects a genuine category ambiguity rather than instrument failure, which is the more analytically consequential finding.

The dominance ratio analysis extends this picture across the full corpus. Human Capital classifications carry the highest textual certainty, a consequence of the stringency of the gating mechanism, which requires a sufficient density of strong category-specific indicators before assignment. The result is a median dominance ratio of approximately 7.3, substantially higher than any other category.

Classifications in Motivational Activation and Capability and Health, where median dominance ratios cluster between 2.0 and 2.3, should be interpreted as indicating the relative prominence of a given logic rather than its exclusive presence. The Other/Unclear and Mixed/Hybrid categories function as analytically honest residuals: they do not represent failures of the instrument but rather its principled recognition of the limits of its own discriminatory capacity.

Secondary classification constitutes the weakest dimension of the instrument's performance. Disagreement arose in a third of sampled cases, with the dominant pattern being the algorithm assigning Work-First as secondary where the LLM assigned Motivational Activation. The most precise interpretation of this pattern is not that Work-First lexis crowds out competing signals, but that Motivational Activation is a logic whose causal mechanism tends to be expressed implicitly in project documentation, through orientation, framing, and attitudinal register rather than through discrete mechanism-specific vocabulary. A dictionary instrument, by design, recovers lexical density rather than discursive spirit, and it is this that the LLM, drawing on contextual and tonal cues, is better positioned to detect. The asymmetry is therefore not incidental but structural, and its consequences are most significant at the secondary level, where Motivational Activation and Social Capital function not as main logics but as hybrid components within more complex intervention designs. Their implicit mode of expression makes them systematically harder to detect as secondary signals. The secondary distribution should accordingly be understood as a conservative lower bound on the occurrence of relational and behavioural mechanisms in the portfolio.

A further property of the instrument bears on how the distributional findings should be read. The dictionary-based approach can provide evidence of the presence of a logic, but cannot straightforwardly be used to infer its absence. Social Capital shows genuinely low prevalence at both primary and secondary levels, consistent with the argument developed in section 7.3 that its mechanisms are structurally difficult to express within European Social Fund documentation. Relational and

network-based effects resist the individual-level output vocabulary that the reporting architecture rewards, and this resistance is reflected in the classification output. Motivational Activation presents a different picture. Despite its modest primary presence, it appears as the second most common secondary logic across the eligible portfolio, suggesting that the instrument does recover motivational mechanisms where they function as complementary rather than dominant signals. The implicit expression argument, therefore, applies most forcefully to Motivational Activation as a primary logic, where it must dominate the signal space to trigger classification, rather than as a secondary signal where its presence alongside a stronger primary logic appears sufficient for detection. Capability and Health presents no comparable detection puzzle: it is present at meaningful levels in both primary and secondary distributions. However, its relatively higher frequency as a secondary logic than its primary share alone would suggest is consistent with the sequential activation logic discussed in section 7.3, where health stabilisation functions as a precondition for rather than the dominant expression of an intervention's theory of change. For all categories, a project in which a logic is not detected should not be interpreted as one from which it is absent; it is one in which it did not produce a signal of sufficient density to meet the classification threshold. The instrument's output is accordingly more precisely characterised as a lower bound on project prevalence than as a complete distributional account. Where a logic is detected, the classification carries genuine evidential weight; where it is not, the absence of a signal should not be read as the absence of the mechanism. This is not a concession about the instrument's reliability but a clarification of its epistemic scope, and it guards against over-interpretation of near-zero counts as evidence of genuine programmatic absence.

Taken together, the evidence supports the conclusion that a theory-driven, dictionary-based instrument can reliably distinguish primary intervention logics in a Swedish-language policy corpus of this type. That conclusion carries an important qualification: classification confidence should be treated as a continuous and variable property rather than a binary outcome, and the dominance and signal-

strength data should be interpreted alongside the distributional results rather than subordinated to them. Within those parameters, the instrument delivers what it was designed to deliver: a transparent, replicable, and theoretically grounded classification of intervention logics at a scale that manual analysis cannot feasibly achieve.

8. Conclusion

8.1 Principal Findings and Their Significance

The European Social Fund's Programme Area 2 portfolio emerges from this analysis as a structured but internally differentiated activation landscape that cannot be adequately described by its modal categories alone. Beneath the surface concentration of Work-First and Human Capital lies substantial variation in what logics involve in practice, how they combine with secondary mechanisms, where they position participants in the labour market trajectory, and how coherently they are expressed in project documentation. The Human Capital share of the portfolio is dominated by language training and formal vocational education rather than the credential-producing investment most directly associated with strong medium-to-long-term returns. Work-First functions as a structuring complement across the portfolio, embedded within projects primarily organised around other mechanisms in ways that suggest a deliberate sequential design logic rather than programmatic incoherence. Functional orientation is not randomly distributed across logics but systematically differentiated. Different theories of change position participants at structurally distinct points along the activation trajectory, with some logics oriented toward stabilisation and others toward direct entry, and Social Capital resisting the sequential model altogether.

The most consequential finding of this study relates to what this configuration reveals about the governance conditions that produce it. The skew towards institutionally legible logics, those that generate measurable outputs, align with

established administrative categories, and match dominant diagnostic frameworks, is not accidental. It reflects a funding structure that shapes which theories of change can be expressed, counted, and rewarded within European Social Fund programming. The marginal status of Social Capital indicates something more specific than mere underrepresentation: a reporting infrastructure that offers no suitable language for relational and network-based mechanisms makes certain theories of change structurally invisible. The portfolio's composition may therefore mirror governance constraints as much as programme choices, and any evaluation ignoring this distinction risks wrongly attributing structural effects to programme design or participant heterogeneity. Practically, a portfolio where logics specifically aimed at addressing functional, attitudinal, and relational barriers occupy relatively marginal positions raises questions about its alignment with the populations it is targeting, those for whom activation readiness cannot be assumed.

The implication is not simply that different interventions should receive more funding, but that the informational infrastructure through which the programme accounts for itself needs to be capable of recognising the theories of change it currently renders inexpressible. A governance system that can only count what it has already decided to measure will continue to produce portfolios skewed toward institutionally legible logics, irrespective of what the evidence on differential effectiveness says about populations being targeted. The instrument developed in this study offers one practical step: by making intervention logic composition visible and replicable at scale, it creates the empirical basis on which programme designers, funders, and evaluators can ask whether the accountability architecture is as well-calibrated to its most disadvantaged participants as it needs to be, and begin to answer the question with evidence rather than assumption.

8.2 What the Validation Evidence Establishes

The evidence on RQ2, whether a theory-driven, dictionary-based instrument can reliably distinguish between distinct intervention logics in a large Swedish-

language policy corpus, supports a positive conclusion within clearly defined limits. Primary classification is robust, with errors structured rather than random, concentrated at theoretically adjacent boundaries where genuine category ambiguity is expected, and detectable through the dominance ratio analysis that accompanies every primary assignment. Classification confidence is not uniform across the typology, but varies predictably in ways that the instrument's own outputs can characterise. Secondary classification is the instrument's weakest dimension, and is best read as a lower bound on mechanism co-occurrence rather than a definitive account. These properties are not incidental limitations but constitutive features of the approach; the price of theoretical transparency and replicability is sensitivity to implicit expression and lexical variation, a trade-off that is deemed appropriate and well-suited to a corpus of this kind. Within those parameters, the instrument is fit for the purpose for which it was designed: not to recover every mechanism present in every document, but to produce theoretically grounded, auditable, and reproducible primary classifications at a scale that manual analysis cannot.

8.3 Contributions

This study makes two related contributions, one substantive and one methodological, each of which conditions the other's validity.

The substantive contribution is the first systematic mapping of intervention logics across the European Social Fund's Programme Area 2 portfolio. By shifting analytical attention from programme effects to programme content, the study produces a structured account of what kinds of activation the programme funds, how those logics are combined and positioned within the activation trajectory, and what the resulting distribution reveals about the governance conditions that produce it. That reorientation matters, since the accumulated evidence on differential programme effectiveness presupposes that one knows what type of intervention is being implemented. This presupposition carries direct consequences for how

effectiveness is assessed, as Card, Kluve, and Weber (2018) and Kluve (2010) demonstrate through their meta-analytic findings on differential returns, and cannot be taken for granted. The typology creates the empirical foundation on which that presupposition can be examined rather than assumed, and in doing so, it opens a line of inquiry that the effects-focused active labour market policy literature has not, to date, systematically pursued. The study contributes to the sociology of classification by demonstrating how institutional infrastructures shape programme content at the portfolio level, extending Bowker and Star's (1999) account of how classificatory infrastructure renders certain practices legible while obscuring others from its origins in scientific and administrative classification to the domain of welfare state activation policy.

The methodological contribution is the demonstration that theory-driven, dictionary-based classification can be applied to Swedish-language programme documents to generate a reliable and interpretable typology of intervention logics at scale. The contribution is not merely technical, but an epistemological precondition under which the substantive contribution carries analytical weight. By maintaining an explicit and auditable link between theoretical categories and textual signals, the instrument makes the classificatory choices visible and contestable in ways that supervised machine learning and unsupervised topic modelling cannot. The validation framework developed alongside the instrument, combining LLM-assisted external validation with dominance ratio analysis across the full corpus, provides a template for assessing classification confidence that is transferable to other corpora and other policy domains.

Together, these contributions respond to a gap that sits at the intersection of substantive and methodological concerns in labour market policy research, the need to complement a well-developed literature on what works with a more systematic understanding of what is being done.

8.4 Limitations

Several limitations merit acknowledgement. The classification instrument is contingent on mechanism-specific vocabulary being present and sufficiently dense in project documentation. Where final reports are written predominantly in governance and process language, specifying partnerships, referral pathways, and institutional arrangements rather than causal logics, the instrument has limited discriminatory purchase. The Other/Unclear category captures some of this residual.

Furthermore, the instrument classifies documentation rather than implementation. Final reports are produced within an institutional framework that shapes how interventions are described, and the discrepancy between educational rhetorics and classified mechanisms identified in the results illustrates that project documentation may not fully represent what is delivered. The findings should accordingly be understood as a characterisation of how European Social Fund Programme Area 2 projects represent their intervention logics in formal documentation rather than as a direct account of the causal mechanisms that operate in implementation.

Classification confidence varies across the typology in ways that are empirically characterised but not eliminable. Human Capital classifications carry the highest textual certainty given the stringency of the gating mechanism. Social Capital results should be read as indicating its relative prominence rather than its exclusive presence. At the secondary level, the distribution represents a conservative lower bound on mechanism co-occurrence rather than a comprehensive account. The functional orientation dimension was internally calibrated rather than externally validated for reasons developed in section 5.3.5; its outputs are accordingly best treated as structurally meaningful patterns within this corpus rather than precise measurements of a theoretically universal construct.

The LLM-assisted validation rests on a sample of 30 projects drawn from a corpus

of 484. This is sufficient to identify systematic patterns in misclassification and to characterise the instrument's behaviour at category boundaries, but it is indicative rather than exhaustive. The accuracy estimate should be interpreted as a well-grounded approximation rather than a precisely calibrated population parameter and is best read in conjunction with the dominance ratio analysis across the full corpus.

Finally, the typology and dictionary are bespoke to the Swedish European Social Fund Programme Area 2 context, constructed through iterative engagement with this specific corpus in Swedish and calibrated to the reporting conventions of the 2014-2020 programme period. This method is designed to be transferable as a procedure and analytical logic, not as a fixed instrument; direct application to other national programmes, funding periods, or policy domains would require systematic re-validation against the target corpus.

8.5 Future research

Several directions follow naturally from this study. The most immediate relates to the connection between portfolio composition and programme effectiveness. The typology developed here makes it possible to ask whether the logics that dominate the portfolio are the same as those the active labour market policy literature identifies as yielding the strongest returns, and for which participant groups. Linking classification output to project-level outcome data would allow that question to be addressed directly, transforming the present descriptive mapping into the foundation for a systematic effectiveness analysis organised around theoretically meaningful categories rather than administrative labels.

A second direction concerns the gap between documented and implemented logic. The present study classifies how projects describe their intervention logic in formal documentation, not what is delivered on the ground. Qualitative case work on a subset of projects, comparing documentary classification against participant and

practitioner accounts of the actual theory of change, would provide a form of implementation validity that the study's scale precluded.

A third direction concerns comparative application. The method is designed to be transferable as a procedure and analytical logic to other European Social Fund funding periods, other national programmes, and potentially other activation corpora. A first step is already underway; the classification instrument is currently being applied to the European Social Fund's Programme Area 1 corpus in collaboration with the Department of Economics at Uppsala University, which will provide the first direct test of transferability and a comparative baseline against which the Programme Area 2 findings can be assessed. More broadly, comparative application would allow the governance argument developed in section 7.3 to be tested against other institutional contexts, examining how far the concentration patterns identified here reflect the specifics of the Swedish programme environment versus the structural incentives of ESF programming more broadly.

8.6 Closing remarks

What this study ultimately demonstrates is that the question of what labour market programmes do, not only what effects they produce, but through what causal logic they are designed to operate, is a tractable empirical question. The instrument developed here shows that systematic, theory-driven mapping of programme content is feasible at scale, and that it reveals things about the governance conditions shaping activation policy that outcome-focused analysis cannot see. Making the intervention logic composition of publicly funded programmes visible is a precondition for asking whether the investment is as well-calibrated as it could be.

9. Bibliography

Arbetsförmedlingen. (2026). *Insatser som leder till jobb? En kunskapsöversikt om effekterna av aktiva arbetsmarknadspolitiska insatser*. Arbetsförmedlingen.

Astbury, B., & Leeuw, F. L. (2010). Unpacking black boxes: Mechanisms and theory building in evaluation. *American Journal of Evaluation*, 31(3), 363–381. <https://doi.org/10.1177/1098214010371972>

Bertoni, E., Fontana, M., Gabrielli, L., Signorelli, S., & Vespe, M. (Eds.). (2023). *Handbook of computational social science for policy*. Springer. <https://doi.org/10.1007/978-3-031-16624-2>

Bonoli, G. (2010). The political economy of active labor-market policy. *Politics & Society*, 38(4), 435–457.

Bonoli, G., & Natali, D. (2012). *The politics of the new welfare state*. Oxford University Press.

Bonvin, J.-M., & Farvaque, N. (2005). What informational basis for assessing job-seekers? Capabilities vs. preferences. *Review of Social Economy*, 63(2), 269–289.

Bowker, G. C., & Star, S. L. (1999). *Sorting things out: Classification and its consequences*. MIT Press.

Calmfors, L. (1994). Active labour market policy and unemployment: A framework for the analysis of crucial design features. *OECD Economic Studies*, 22, 7–47.

Card, D., Kluve, J., & Weber, A. (2018). What works? A meta-analysis of recent active labor market program evaluations. *Journal of the European Economic Association*, 16(3), 894–931.

- Chetty, R. (2015). Behavioral economics and public policy: A pragmatic perspective. *American Economic Review*, 105(5), 1–33.
- Deming, D., & Silliman, M. (2025). Skills and human capital in the labor market. In C. Dustmann & T. Lemieux (Eds.), *Handbook of labor economics* (Vol. 6, pp. 115–157). Elsevier. <https://doi.org/10.1016/bs.heslab.2025.07.001>
- Douglas, M. (1986). *How institutions think*. Syracuse University Press.
- Durkheim, É., & Mauss, M. (1963). *Primitive classification* (R. Needham, Trans.). Cohen & West. (Original work published 1903)
- Engdahl, M., Forslund, A., & Vikman, U. (2023). *(Kommunala) insatser för att underlätta arbetsmarknadsinträdet för flyktingar och deras anhöriga* (IFAU Rapport 2023:8). IFAU.
- Forslund, A. (2018). *Subventionerade anställningar — avvägningar och empirisk evidens* (IFAU Rapport 2018:14). IFAU.
- Granovetter, M. (1983). The strength of weak ties: A network theory revisited. *Sociological Theory*, 1, 201–233.
- Grimmer, J., & Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis*, 21(3), 267–297.
- Kluve, J. (2010). The effectiveness of European active labor market programs. *Labour Economics*, 17(6), 904–918.
- Pawson, R., & Tilley, N. (1997). *Realistic evaluation*. SAGE Publications.

Putnam, R. D. (2000). *Bowling alone: The collapse and revival of American community*. Simon & Schuster.

Ramboll Management Consulting. (2025). *Effektutvärdering av arbete och studier: Metodrapport* (Delrapport, september 2025). Ramboll Management Consulting / Svenska ESF-rådet.

Riksrevisionen. (2017). *Förberedande och orienterande utbildning* (RiR 2017:20). Riksrevisionen.

Sainsbury, R., Coleman-Fountain, E., & Trezzini, B. (2017). How to enhance active citizenship for persons with disabilities in Europe through labour market participation: European and national perspectives. In R. Halvorsen, B. Hvinden, J. Bickenbach, D. Ferri, & A. M. Guillén Rodríguez (Eds.), *The changing disability policy system: Active citizenship and disability in Europe volume 1* (pp. 90–107). Routledge. <https://doi.org/10.4324/9781315623931-6>

Sunstein, C. R., & Thaler, R. H. (2003). Libertarian paternalism is not an oxymoron. *University of Chicago Law Review*, 70(4), 1159–1202.

Svenska ESF-rådet. (n.d.). *Europeiska socialfonden i Sverige: Program för investering i tillväxt och sysselsättning 2014–2020*. Svenska ESF-rådet. https://www.esf.se/app/uploads/2021/01/ESF_popversion_sidor_LOW3.compressed-1.pdf

Svenska ESF-rådet. (2022). *Nationellt socialfondsprogram för investering för tillväxt och sysselsättning 2014–2020* (Version 10.0). Svenska ESF-rådet. https://www.esf.se/app/uploads/2024/11/2024-00670-2-Operativt-program-2014-2020-10.0-1413525_1_1.pdf

Sweco. (2021). *Effektutvärdering av socialfondsprogrammet 2014–2020* (Diarienummer 2020/00531). Svenska ESF-rådet.

Weiss, C. H. (1997). Theory-based evaluation: Past, present, and future. *New Directions for Evaluation*, 1997(76), 41–55. <https://doi.org/10.1002/ev.1086>

10. Appendices

10.1 Appendix A

The following tables present the complete keyword dictionaries applied in the classification procedure. All keywords are in Swedish, reflecting the language of the European Social Fund Programme Area 2 final report corpus. Table A1 presents the primary intervention logic dictionary. For Human Capital, keywords are divided into three weighted tiers: strong, medium and weak indicators. Generic phrases excluded prior to scoring are listed separately. For all other categories, keywords are unweighted and counted as raw frequencies normalised per 1,000 words. Several terms appear across multiple dictionaries by design, as the primary logics and functional orientation dimensions are analytically independent.

Table A1. Primary Intervention Logic Dictionary

Category	Tier / Type	Keywords (Swedish)
Human Capital	Strong (weight: 3)	yrkesutbildning, yrkesutbildningen, yrkesutbildningar, yrkesutbildningarna, utbildningsprogram, utbildningsprogrammet, utbildningsprogrammen, certifiering, certifieringen, certifieringar, certifieringarna, certifikat, certifikatet, certifikaten, diplom, diplommet, diplomerna, validering, valideringen, valideringar, valideringarna, reell kompetens, reella kompetenser, komvux, examen
	Medium (weight: 2)	utbildningsinsats, utbildningsinsatsen, utbildningsinsatser, utbildningsinsatserna, språkträning, språkträningen, studier,

Category	Tier / Type	Keywords (Swedish)
		studierna, sfi, svenska för invandrare, språkutbildning, språkutbildningen, språkutbildningar, språkkurs, språkkursen, språkkurser, språkkurserna, vuxenutbildning, vuxenutbildningen, vuxenutbildningar, yrkessvenska, yrkessvenskan, orientering, orienteringen, orienteringar, preparandutbildning, preparandutbildningen, introduktionsskurs, introduktionsskursen, introduktionsskurser
	Weak (weight: 1)	utbildning, utbildningen, utbildningar, utbildningarna, kurs, kursen, kurser, kurserna, kompetensutveckling, kompetensutvecklingen, kompetenshöjande, kunskap, kunskapen, kunskaper, kunskaperna, kunskapsutveckling, kunskapsutvecklingen, yrkeskompetens, yrkeskompetensen, formell kompetens, formella kompetenser
	Generic exclusions	öka kompetensen, höja kompetensen, kompetens för framtiden, kompetensbehov, kompetensförsörjning, lära, lärande, kompetens, kompetensen, kompetenserna
Work-First	Unweighted	praktik, praktiken, arbetspraktik, arbetspraktiken, praktikplats, praktikplatsen, praktikplatser, praktikplatserna, arbetsplatsförlagt, arbetsplatsförlagd, arbetsplatsförlagda, anställning, anställningen, anställningar, anställningarna, lönebidrag, lönebidraget, lönebidragen, nystartsjobb, nystartsjobbet, nystartsjobben, instegsjobb, instegsjobbet, instegsjobben, traineejobb, traineejobbet, subventionerad anställning, subventionerade anställningar, jobbmatchning, jobbmatchningen, matchningstjänst, matchningstjänsten, matchningstjänster, snabbspår, snabbspåret, snabbspåren, lönesstöd, lönesstödet, anställningsstöd, anställningsstödet, jobbsökning, jobbsökningen, jobbsökande, rekryteringsstöd, rekryteringsstödet
Capability and Health	Unweighted	arbetsförmåga, arbetsförmågan, funktionsnedsättning, funktionsnedsättningen, funktionsnedsättningar, funktionsnedsättningarna, rehabilitering, rehabiliteringen, rehabiliteringar, psykisk hälsa, psykiska hälsan, psykisk ohälsa, psykiska ohälsan, stress, stressen, utmattning, utmattningen, utmattningssyndrom, återhämtning, återhämtningen, hälsofrämjande, sjukskrivning,

Category	Tier / Type	Keywords (Swedish)
		sjukskrivningen, sjukskrivningar, sjukskrivningarna, sjukersättning, sjukersättningen, funktionshinder, funktionshindret, funktionshindren, beroendeproblematik, beroendeproblematiken, neuropsykiatrisk, neuropsykiatriska, neuropsykiatriskt, missbruk, missbruket, samsjuklighet, samsjukligheten, arbetsförmågebedömning, arbetsförmågebedömningen, arbetsförmågebedömningar
Social Capital	Unweighted	nätverkande, nätverkandet, mentorskap, mentorskapet, mentor, mentorer, föreningsliv, föreningslivet, kontaktnät, kontaktnätet, professionellt nätverk, yrkesnätverk, yrkesnätverket, brobyggande, brobyggandet, ideellt arbete, ideella arbetet, social tillhörighet, social tillhörigheten, kamratstöd, kamratstödet, fadder, faddern, faddrar, faddrarna, ambassadör, ambassadören, ambassadörer, mötesplats, mötesplatsen, mötesplatser, mötesplatserna, branschkontakter, branschkontakterna, branschkontakt, branschkontakten
Motivational Activation	Unweighted	motivation, motivationen, motivationer, motiverande, motiverar, motivera, coachning, coachningen, coaching, empowerment, resiliens, resiliensen, självförtroende, självförtroendet, självkänsla, självkänslan, attitydförändring, attitydförändringar, drivkraft, drivkraften, drivkrafter, beteendeförändring, beteendeförändringen, beteendeförändringar, stödsamtal, stödsamtalet, stödsamtalen, framtidstro, framtidstron, personlig utveckling, personliga utvecklingen, meningsfull, meningsfulla, meningsfullt, självanalys, självanalysen, mindset, framtidsplanering, framtidsplaneringen, livsplan, livsplanen, livsplaner, dröm, drömmen, drömmar, drömmarna, självbild, självbilden, ökad tillit, ökade tillit, tillit, tilliten

Each functional orientation flag is triggered when both a minimum raw keyword count and a minimum normalised rate per 1,000 words are met. Where more than one flag is triggered simultaneously, the project is assigned Mixed/multiple. Where no flag is triggered, the project is assigned Other.

Table A2. Functional Orientation Dictionary

Functional Orientation	Detection Thresholds	Keywords (Swedish)
Assessment	$n \geq 2$; rate ≥ 0.20	kartläggningsverktyg, kartläggningsverktyget, kompetenskartläggning, kompetenskartläggningen, kompetenskartläggningar, kompetensprofil, kompetensprofilen, kompetensprofiler, arbetsförmågebedömning, arbetsförmågebedömningen, arbetsförmågebedömningar, nivåtest, nivåtestet, nivåtester, screening, screeningen
Stabilising	$n \geq 3$; rate ≥ 0.30	rehabilitering, rehabiliteringen, rehabiliteringar, psykisk ohälsa, psykiska ohälsan, psykisk hälsa, psykiska hälsan, välbefinnande, välbefinnandet, stress, stressen, utmattning, utmattningen, utmattningssyndrom, återhämtning, återhämtningen, arbetslivsinriktad rehabilitering, arbetslivsinriktade rehabilitering
Direct Labour Market Entry	$n \geq 3$; rate ≥ 0.35	praktik, praktiken, anställning, anställningen, anställningar, anställningarna, matchning, matchningen, lönebidrag, lönebidraget, lönebidragen, nystartsjobb, nystartsjobbet, nystartsjobben, arbetspraktik, arbetspraktiken, instegsjobb, instegsjobbet, instegsjobben, subventionerad anställning, subventionerade anställningar, ingångsjobb, ingångsjobbet, ingångsjobben
Preparatory	$n \geq 2$; rate ≥ 0.20	intervjuträning, intervjuträningen, jobbklubb, jobbklubben, jobbklubbar, jobbklubbarna, intervjucoachning, intervjucoachningen, jobbsök, jobbsökning, jobbsökningen, arbetsmarknadskunskap, arbetsmarknadskunskapen, introduktionsutbildning, introduktionsutbildningen, introduktionsutbildningar, förberedande utbildning, förberedande utbildningen, karriärvägledning, karriärvägledningen, studiecirkel, studiecirkeln, studiecirkelar, studiecirkelarna, förberedande insats, förberedande insatser

Subtypes are assigned on a frequency basis; all four subtype dictionaries are scored simultaneously, and the subtype with the highest total keyword count is assigned. Raw keyword counts are used rather than normalised rates, as subtype assignment is a within-document comparison in which document length affects all subtypes equally. HC5 is a residual category assigned when no subtype-specific keywords are detected across any of the four dictionaries.

Table A3: Human Capital Subtype Classification Keywords

Subtype	Label	Keywords (Swedish)
HC1	Credential	certifiering, certifieringen, certifieringar, certifieringarna, certifikat, certifikatet, certifikaten, diplom, diplomet, diplomen, behörighet, behörigheten, behörigheter, behörigheterna, yrkeslegitimation, yrkeslegitimationen, yrkeslicens, yrkeslicensen
HC2	Language	sfi, språkutbildning, språkutbildningen, språkutbildningar, språkträning, språkträningen, svenska för invandrare, modersmål, modersmålet, språkstöd, språkstödet, svenska som andraspråk, sva, yrkessvenska, yrkessvenskan
HC3	Validation	validering, valideringen, valideringar, valideringarna, reell kompetens, reella kompetenser, kompetensbevis, kompetensbeviset, yrkesbedömning, yrkesbedömningen, yrkesbedömningar, kompetensbedömning, kompetensbedömningen, kompetensbedömningar, recognition, kompetensanalys, kompetensanalysen, kompetensanalyser, utländska meriter, utländsk merit, branschvalidering, branschvalideringen, branschvalideringar
HC4	Formal / Vocational	yrkesutbildning, yrkesutbildningen, yrkesutbildningar, yrkesutbildningarna, utbildningsprogram, utbildningsprogrammen, vuxenutbildning, vuxenutbildningen, vuxenutbildningar, yrkeskurs, yrkeskursen, yrkeskurser, yrkesförberedande, komvux, lärling, lärlingen, lärlingar, lärlingarna
HC5	Other Structured Training	Assigned when project is classified as Human Capital but no HC1-HC4 keywords are detected

10.2 Appendix B

This appendix documents the complete set of decision rules and numerical thresholds applied in the classification procedure. All values are fixed and not adjusted following the LLM-assisted validation reported in section 5.4. Rules are applied sequentially in the order presented. The dominance ratio is defined as the ratio of the highest-scoring category score to the second-highest. The secondary share is defined as the ratio of the second-highest score to the highest.

Table B1. Corpus Inclusion Thresholds

Parameter	Threshold	Rationale
Minimum character count	$\geq 1,000$ characters	Ensures sufficient text for extraction; documents below this threshold typically indicate extraction failure or near-empty files
Minimum word count	> 500 words	Ensures sufficient textual signal for reliable keyword scoring across all categories

Both gates are evaluated individually. A project meeting either Gate A or Gate B is eligible for Human Capital classification. HC_rate is the weighted score (3x strong + 2x medium + 1x weak hits) normalised per 1,000 words, calculated after removing generic exclusion phrases.

Table B2. Human Capital Gating Conditions

Rule	Condition	Notes
Gate A	HC_strong_n ≥ 3 AND HC_rate ≥ 12	Requires at least three strong indicators with a minimum weighted rate. Primary gate for most HC classifications.

Gate B	$HC_rate \geq 15$	High overall weighted rate sufficient on its own, even without three strong indicators. Captures dense but varied HC vocabulary.
Outcome if neither gate met	Human Capital set to $-\infty$	Project cannot be classified as Human Capital regardless of raw score. Eligible for all other categories.

Applied after gating. The dominance ratio compares the top-scoring to the second-highest-scoring category. A ratio below 1.3 indicates insufficient separation between the two leading categories.

Table B3. Primary Classification Decision Rules

Decision Rule	Condition	Assigned Category
No detectable signal	$top1 = -\infty$	Other/Unclear
Signal below minimum threshold	$top1 < 0.15$	Other/Unclear
Competing categories too close	$dominance_ratio < 1.3$	Mixed/Hybrid
Clear dominant category	$top1 \geq 0.15$ AND $dominance_ratio \geq 1.3$	Highest-scoring eligible category

Applied only to projects with a substantive primary classification. Secondary classification is restricted to projects not assigned Mixed/Hybrid or Other/Unclear as their primary type.

Table B4. Secondary Classification Rules

Rule	Condition	Notes
Eligible projects	Primary type not Mixed/Hybrid or Other/Unclear	387 of 484 projects eligible for secondary classification

Secondary share threshold	$\text{secondary_share} \geq 0.20$	Second-ranked category must score at least 20% of the top-scoring category. Conservative rule to prevent spurious secondary assignments.
Human Capital as secondary	Ineligible	HC gating applies to secondary classification. HC cannot be assigned as secondary logic.
Outcome if threshold not met	$\text{secondary_type} = \text{NA}$	Project assigned no secondary logic. Does not imply the absence of secondary mechanisms, only that no second logic crossed the detection threshold.

Applied only to projects classified as primary Human Capital. Subtypes are assigned based on the relative frequency of keyword hits across all four subtype dictionaries. Each subtype dictionary is scored simultaneously, and the subtype with the highest total keyword count is assigned. Since subtype assignment is a within-document comparison, identifying which subtype has the highest keyword frequency relative to others in the same text, normalisation by document length would not alter the relative ranking and is therefore not applied.

Table B5. Human Capital Subtype Assignment

Priority	Subtype	Trigger Condition
Frequency-based	HC1 Credential	Assigned when HC1 keywords count is highest
Frequency-based	HC2 Language	Assigned when HC2 keywords count is highest
Frequency-based	HC3 Validation	Assigned when HC3 keywords count is highest
Frequency-based	HC4 Formal/Vocational	Assigned when HC4 keywords count is highest
Frequency-based	HC5 Other Training	Assigned when total keyword count across HC1-HC4 is zero

Each functional flag requires both a minimum raw keywords count and a minimum normalised rate. When no flag is triggered, the project is assigned Other. Where more than one flag is triggered simultaneously the project is assigned with Mixed/multiple.

Table B6. Functional Orientation Flags and Thresholds

Function	Min. Count	Min. Rate (per 1,000 words)	Multi-flag outcome
Assessment	$n \geq 2$	rate ≥ 0.20	If more than one flag triggered: Mixed/Multiple
Stabilising	$n \geq 3$	rate ≥ 0.30	
Direct Labour Market Entry	$n \geq 3$	rate ≥ 0.35	
Preparatory	$n \geq 2$	rate ≥ 0.20	
Other	—	—	Assigned when no flag is triggered

10.3 Appendix C

R script

```
# ESF TYPOLOGY CLASSIFICATION
# Louise MacRae - SSDA Thesis project
options(stringsAsFactors = FALSE)
set.seed(123)

packages <- c("tidyverse", "stringr", "purrr", "pdftools", "writexl", "readr")
installed <- packages %in% rownames(installed.packages())
if (any(!installed)) install.packages(packages[!installed])
```

```
library(tidyverse)
```

```
library(stringr)
```

```
library(purrr)
```

```
library(pdftools)
```

```
library(writexl)
```

```
library(readr)
```

```
BASE_DIR <- "/Users/louisemacrae/ESF"
```

```
# Output folder
```

```
OUT_DIR <- file.path(BASE_DIR, "outputs_typology_PO2")
```

```
dir.create(OUT_DIR, showWarnings = FALSE, recursive = TRUE)
```

```
# Dictionaries
```

```
hc_strong <- c(
```

```
"yrkesutbildning", "yrkesutbildningen", "yrkesutbildningar", "yrkesutbildningarna",
```

```
"utbildningsprogram", "utbildningsprogrammet", "utbildningsprogrammen",
```

```
"certifiering", "certifieringen", "certifieringar", "certifieringarna",
```

```
"certifikat", "certifikatet", "certifikaten",
```

```
"diplom", "diplomet", "diplomen",
```

```
"validering", "valideringen", "valideringar", "valideringarna",
```

```
"reell kompetens", "reella kompetenser",
```

```
"komvux", "examen"
```

```
)
```

```
hc_medium <- c(
```

"utbildningsinsats", "utbildningsinsatsen", "utbildningsinsatser", "utbildningsinsatserna",

"språkträning", "språkträningen",

"studier", "studierna",

"sfi", "svenska för invandrare",

"språkutbildning", "språkutbildningen", "språkutbildningar",

"språkkurs", "språkkursen", "språkkurser", "språkkurserna",

"vuxenutbildning", "vuxenutbildningen", "vuxenutbildningar",

"yrkessvenska", "yrkessvenskan",

"orientering", "orienteringen", "orienteringar",

"preparandutbildning", "preparandutbildningen",

"introduktionskurs", "introduktionskursen", "introduktionskurser"

)

hc_weak <- c(

"utbildning", "utbildningen", "utbildningar", "utbildningarna",

"kurs", "kursen", "kurser", "kurserna",

"kompetensutveckling", "kompetensutvecklingen",

"kompetenshöjande",

"kunskap", "kunskapen", "kunskaper", "kunskaperna",

"kunskapsutveckling", "kunskapsutvecklingen",

"yrkeskompetens", "yrkeskompetensen",

"formell kompetens", "formella kompetenser"

)

hc_generic_exclude <- c(

"öka kompetensen", "höja kompetensen", "kompetens för framtiden",

"kompetensbehov", "kompetensförsörjning", "lära", "lärande", "kompetens",

"kompetensen", "kompetenserna"

)

```

wf_terms <- c(
  "praktik", "praktiken",
  "arbetspraktik", "arbetspraktiken",
  "praktikplats", "praktikplatsen", "praktikplatser", "praktikplatserna",
  "arbetsplatsförlagt", "arbetsplatsförlagd", "arbetsplatsförlagda",
  "anställning", "anställningen", "anställningar", "anställningarna",
  "lönebidrag", "lönebidraget", "lönebidragen",
  "nystartsjobb", "nystartsjobbet", "nystartsjobben",
  "instegsjobb", "instegsjobbet", "instegsjobben",
  "traineejobb", "traineejobbet",
  "subventionerad anställning", "subventionerade anställningar",
  "jobbmatchning", "jobbmatchningen",
  "matchningstjänst", "matchningstjänsten", "matchningstjänster",
  "snabbspår", "snabbspåret", "snabbspåren",
  "lönestöd", "lönestödet",
  "anställningsstöd", "anställningsstödet",
  "jobsökning", "jobsökningen",
  "jobsökande",
  "rekryteringsstöd", "rekryteringsstödet"
)

ch_terms <- c(
  "arbetsförmåga", "arbetsförmågan",

  "funktionsnedsättning", "funktionsnedsättningen", "funktionsnedsättningar", "funktionsnedsättningarna",
  "rehabilitering", "rehabiliteringen", "rehabiliteringar",
  "psykisk hälsa", "psykiska hälsan",
  "psykisk ohälsa", "psykiska ohälsan",
  "stress", "stressen",

```

"utmattning", "utmattningen", "utmattningssyndrom",
 "återhämtning", "återhämtningen",
 "hälsofrämjande",
 "sjukskrivning", "sjukskrivningen", "sjukskrivningar", "sjukskrivningarna",
 "sjukersättning", "sjukersättningen",
 "funktionshinder", "funktionshindret", "funktionshindren",
 "beroendeproblematik", "beroendeproblematiken",
 "neuropsykiatrisk", "neuropsykiatriska", "neuropsykiatriskt",
 "missbruk", "missbruket",
 "samsjuklighet", "samsjukligheten",

 "arbetsförmågebedömning", "arbetsförmågebedömningen", "arbetsförmågebedömn
 ingar"
)

```

sc_terms <- c(
  "nätverkande", "nätverkandet",
  "mentorskap", "mentorskapet",
  "mentor", "mentorer",
  "föreningsliv", "föreningslivet",
  "kontaktnät", "kontaktnätet",
  "professionellt nätverk",
  "yrkesnätverk", "yrkesnätverket",
  "brobyggande", "brobyggandet",
  "ideellt arbete", "ideella arbetet",
  "social tillhörighet", "social tillhörigheten",
  "kamratstöd", "kamratstödet",
  "fadder", "faddern", "faddrar", "faddrarna",
  "ambassadör", "ambassadören", "ambassadörer",
  "mötesplats", "mötesplatsen", "mötesplatser", "mötesplatserna",
  "branschkontakter", "branschkontakterna", "branschkontakt", "branschkontakten"
)

```

)

```
mo_terms <- c(
  "motivation", "motivationen", "motivationer",
  "motiverande", "motiverar", "motivera",
  "coachning", "coachningen",
  "coaching",
  "empowerment",
  "resiliens", "resiliensen",
  "självförtroende", "självförtroendet",
  "självkänsla", "självkänslan",
  "attitydförändring", "attitydförändringar",
  "drivkraft", "drivkraften", "drivkrafter",
  "beteendeförändring", "beteendeförändringen", "beteendeförändringar",
  "stödsamtal", "stödsamtalet", "stödsamtalen",
  "framtidstro", "framtidstron",
  "personlig utveckling", "personliga utvecklingen",
  "meningsfull", "meningsfulla", "meningsfullt",
  "självanalys", "självanalysen",
  "mindset",
  "framtidsplanering", "framtidsplaneringen",
  "livsplan", "livsplanen", "livsplaner",
  "dröm", "drömmen", "drömmar", "drömmarna",
  "självbild", "självbilden",
  "ökad tillit", "ökade tillit", "tillit", "tilliten"
)
```

Keyword counter and scoring functions

```
keyword_count <- function(text, terms) {
  text <- tolower(text)
```

```

terms <- tolower(terms)
terms_esc <- stringr::str_replace_all(terms, "([\\W])", "\\\\\\\1")
pattern <- paste0("\\b(", paste(terms_esc, collapse = "|"), ")\\b")
out <- stringr::str_count(text, stringr::regex(pattern, ignore_case = TRUE))
out[is.na(text)] <- NA_integer_
as.integer(out)
}

# Alias for vectorised use
keyword_count_vec <- keyword_count

hc_weighted_score <- function(text) {
  t <- tolower(text)
  for (ph in hc_generic_exclude) {
    t <- str_replace_all(t, fixed(tolower(ph)), " ")
  }
  strong_n <- keyword_count_vec(t, hc_strong)
  med_n <- keyword_count_vec(t, hc_medium)
  weak_n <- keyword_count_vec(t, hc_weak)
  score <- 3 * strong_n + 2 * med_n + 1 * weak_n
  list(score = score, strong_n = strong_n, med_n = med_n, weak_n = weak_n)
}

# Load projects
pdf_dir <- file.path(BASE_DIR, "Uppsats/ESF data")

pdf_files <- sort(list.files(
  pdf_dir,
  pattern = "\\..pdf$",
  full.names = TRUE,

```

```

ignore.case = TRUE
))

df_raw <- tibble(
  id = seq_along(pdf_files),
  filename = basename(pdf_files),
  path = pdf_files
) %>%
mutate(
  text = map_chr(path, ~{
    pages <- tryCatch(pdf_text(.x), error = function(e) NULL)
    if (is.null(pages)) return(NA_character_)
    paste(pages, collapse = " ") %>%
      str_replace_all("\\s+", " ") %>%
      str_trim()
  }),
  char_count = if_else(is.na(text), NA_integer_, nchar(text)),
  word_count = str_count(text, "\\w+")
)

# Exclusion log
excluded <- df_raw %>%
mutate(
  exclude_reason = case_when(
    is.na(text) ~ "pdf_text failed / NA",
    char_count < 1000 ~ "char_count < 1000",
    word_count <= 500 ~ "word_count <= 500",
    TRUE ~ NA_character_
  )
) %>%
filter(!is.na(exclude_reason)) %>%

```



```

select(id, filename, path, char_count, word_count, exclude_reason)

write_csv(excluded, file.path(OUT_DIR, "excluded_files.csv"))

# Keep only valid documents
df <- df_raw %>%
  filter(!is.na(text), char_count >= 1000, word_count > 500)

cat("\n--- SAMPLE SUMMARY ---\n")
cat("Total PDFs: ", nrow(df_raw), "\n", sep = "")
cat("Excluded: ", nrow(excluded), "\n", sep = "")
cat("Included: ", nrow(df), "\n", sep = "")

# Keyword counts and normalised rates
df <- df %>%
  mutate(
    hc_tmp = map(text, hc_weighted_score),
    HC_score = map_dbl(hc_tmp, "score"),
    HC_strong_n = map_dbl(hc_tmp, "strong_n"),
    HC_med_n = map_dbl(hc_tmp, "med_n"),
    HC_weak_n = map_dbl(hc_tmp, "weak_n"),

    WF_n = keyword_count(text, wf_terms),
    CH_n = keyword_count(text, ch_terms),
    SC_n = keyword_count(text, sc_terms),
    MO_n = keyword_count(text, mo_terms)
  ) %>%
  select(-hc_tmp) %>%
  mutate(
    HC_rate = HC_score / word_count * 1000,

```

```

    WF_rate = WF_n / word_count * 1000,
    CH_rate = CH_n / word_count * 1000,
    SC_rate = SC_n / word_count * 1000,
    MO_rate = MO_n / word_count * 1000
  )

# Primary logic classification
type_matrix <- df %>%
  select(HC_rate, WF_rate, CH_rate, SC_rate, MO_rate) %>%
  as.matrix()

colnames(type_matrix) <- c(
  "Human Capital",
  "Work-First",
  "Capability/Health",
  "Social Capital",
  "Motivational"
)

# Gating: Human Capital requires minimum signal density
hc_gate <-
  (df$HC_strong_n >= 3 & df$HC_rate >= 12) |
  (df$HC_rate >= 15)

df <- df %>% mutate(hc_gate = hc_gate)

type_matrix_gated <- type_matrix
type_matrix_gated[!hc_gate, "Human Capital"] <- -Inf

# Primary and secondary classification

```

```

primary_idx <- max.col(type_matrix_gated, ties.method = "first")
df$primary_type_raw <- colnames(type_matrix_gated)[primary_idx]

type_matrix_secondary <- type_matrix
type_matrix_secondary[!hc_gate, "Human Capital"] <- -Inf

secondary_idx <- sapply(seq_len(nrow(type_matrix_secondary)), function(i) {
  x <- type_matrix_secondary[i, ]
  x[df$primary_type_raw[i]] <- -Inf
  if (all(is.infinite(x))) return(NA_integer_)
  which.max(x)
})

df$secondary_type_raw <- ifelse(
  is.na(secondary_idx),
  NA_character_,
  colnames(type_matrix_secondary)[secondary_idx]
)

# Dominance ratio
top1 <- apply(type_matrix_gated, 1, max)
top2 <- apply(type_matrix_gated, 1, function(x) {
  vals <- sort(x[is.finite(x)], decreasing = TRUE)
  if (length(vals) < 2) return(0)
  vals[2]
})

dominance_ratio <- ifelse(top2 > 0, top1 / top2, Inf)
dominance_ratio_finite <- ifelse(is.finite(dominance_ratio), dominance_ratio,
NA_real_)
secondary_share <- ifelse(top1 > 0, top2 / top1, NA_real_)

```

```

min_signal <- 0.15
mixed_ratio <- 1.3

df <- df %>%
  mutate(
    top1 = top1,
    top2 = top2,
    dominance_ratio = dominance_ratio,
    dominance_ratio_finite = dominance_ratio_finite,
    secondary_share = secondary_share,
    primary_type = case_when(
      is.infinite(top1) ~ "Other/Unclear",
      top1 < min_signal ~ "Other/Unclear",
      (top2 > 0 & dominance_ratio < mixed_ratio) ~ "Mixed/Hybrid",
      TRUE ~ primary_type_raw
    ),
    secondary_type = case_when(
      primary_type %in% c("Mixed/Hybrid", "Other/Unclear") ~ NA_character_,
      is.na(secondary_type_raw) ~ NA_character_,
      secondary_share < 0.20 ~ NA_character_,
      TRUE ~ secondary_type_raw
    )
  )

# HC subtypes
df <- df %>%
  mutate(
    hc1_n = keyword_count_vec(text, c(
      "certifiering", "certifieringen", "certifieringar", "certifieringarna",

```

```

"certifikat","certifikatet","certifikaten",
"diplom","diplomet","diplomen",
"behörighet","behörigheten","behörigheter","behörigheterna",
"yrkeslegitimation","yrkeslegitimationen",
"yrkeslicens","yrkeslicensen"
)),
hc2_n = keyword_count_vec(text, c(
  "sfi",
  "språkutbildning","språkutbildningen","språkutbildningar",
  "språkträning","språkträningen",
  "svenska för invandrare",
  "modersmål","modersmålet",
  "språkstöd","språkstödet",
  "svenska som andraspråk","sva",
  "yrkessvenska","yrkessvenskan"
)),
hc3_n = keyword_count_vec(text, c(
  "validering","valideringen","valideringar","valideringarna",
  "reell kompetens","reella kompetenser",
  "kompetensbevis","kompetensbeviset",
  "yrkesbedömning","yrkesbedömningen","yrkesbedömningar",
  "kompetensbedömning","kompetensbedömningen","kompetensbedömningar",
  "recognition",
  "kompetensanalys","kompetensanalysen","kompetensanalyser",
  "utländska meriter","utländsk merit",
  "branschvalidering","branschvalideringen","branschvalideringar"
)),
hc4_n = keyword_count_vec(text, c(
  "yrkesutbildning","yrkesutbildningen","yrkesutbildningar","yrkesutbildningarna",
  "utbildningsprogram","utbildningsprogrammen",

```

```

"vuxenutbildning", "vuxenutbildningen", "vuxenutbildningar",
"yrkeskurs", "yrkeskursen", "yrkeskurser",
"yrkesförberedande",
"komvux",
"lärling", "lärlingen", "lärlingar", "lärlingarna"
)),

```

```

HC_subtype = case_when(
  primary_type != "Human Capital" ~ NA_character_,
  (hc1_n + hc2_n + hc3_n + hc4_n) == 0 ~ "HC5 Other training",
  hc1_n >= hc2_n & hc1_n >= hc3_n & hc1_n >= hc4_n ~ "HC1 Credential",
  hc2_n >= hc1_n & hc2_n >= hc3_n & hc2_n >= hc4_n ~ "HC2 Language",
  hc3_n >= hc1_n & hc3_n >= hc2_n & hc3_n >= hc4_n ~ "HC3 Validation",
  hc4_n >= hc1_n & hc4_n >= hc2_n & hc4_n >= hc3_n ~ "HC4
Formal/Vocational",
  TRUE ~ "HC5 Other training"
)
)

```

```

# Functional orientation

```

```

assessment_terms <- c(
  "kartläggningsverktyg", "kartläggningsverktyget",
  "kompetenskartläggning", "kompetenskartläggningen", "kompetenskartläggningar"
,
  "kompetensprofil", "kompetensprofilen", "kompetensprofiler",
  "arbetsförmågebedömning", "arbetsförmågebedömningen", "arbetsförmågebedömn
ingar",
  "nivåtest", "nivåtestet", "nivåtester",

```

```

"screening","screeningen"
)

stabilising_terms <- c(
  "rehabilitering","rehabiliteringen","rehabiliteringar",
  "psykisk ohälsa","psykiska ohälsan",
  "psykisk hälsa","psykiska hälsan",
  "välbefinnande","välbefinnandet",
  "stress","stressen",
  "utmattning","utmattningen","utmattningssyndrom",
  "återhämtning","återhämtningen",
  "arbetslivsinriktad rehabilitering","arbetslivsinriktade rehabilitering"
)

direct_terms <- c(
  "praktik","praktiken",
  "anställning","anställningen","anställningar","anställningarna",
  "matchning","matchningen",
  "lönebidrag","lönebidraget","lönebidragen",
  "nystartsjobb","nystartsjobbet","nystartsjobben",
  "arbetspraktik","arbetspraktiken",
  "instegsjobb","instegsjobbet","instegsjobben",
  "subventionerad anställning","subventionerade anställningar",
  "ingångsjobb","ingångsjobbet","ingångsjobben"
)

preparatory_terms <- c(
  "intervjuträning","intervjuträningen",
  "jobbklubb","jobbklubben","jobbklubbar","jobbklubbarna",
  "intervjucoachning","intervjucoachningen",
  "jobbsök","jobbsökning","jobbsökningen",

```

```

"arbetsmarknadskunskap", "arbetsmarknadskunskapen",

"introduktionsutbildning", "introduktionsutbildningen", "introduktionsutbildningar"
,
"förberedande utbildning", "förberedande utbildningen",
"karriärvägledning", "karriärvägledningen",
"studiecirkel", "studiecirkeln", "studiecirklar", "studiecirkarna",
"förberedande insats", "förberedande insatser"
)

```

```

df <- df %>%
  mutate(
    assessment_n = keyword_count(text, assessment_terms),
    stabilising_n = keyword_count(text, stabilising_terms),
    direct_n     = keyword_count(text, direct_terms),
    preparatory_n = keyword_count(text, preparatory_terms),

    assessment_rate = assessment_n / word_count * 1000,
    stabilising_rate = stabilising_n / word_count * 1000,
    direct_rate     = direct_n     / word_count * 1000,
    preparatory_rate = preparatory_n / word_count * 1000,

    assessment_flag = assessment_n >= 2 & assessment_rate >= 0.20,
    stabilising_flag = stabilising_n >= 3 & stabilising_rate >= 0.30,
    direct_flag     = direct_n     >= 3 & direct_rate     >= 0.35,
    preparatory_flag = preparatory_n >= 2 & preparatory_rate >= 0.20,

    n_flags = assessment_flag + stabilising_flag + direct_flag + preparatory_flag,

    intervention_function = case_when(
      n_flags == 0 ~ "Other",

```



```

    n_flags > 1 ~ "Mixed/Multiple",
    assessment_flag ~ "Assessment",
    stabilising_flag ~ "Stabilising",
    direct_flag ~ "Direct labour market entry",
    preparatory_flag ~ "Preparatory",
    TRUE ~ NA_character_
  )
)

# Multi-label flags
max_hits <- apply(type_matrix, 1, max)
threshold <- 0.3 * max_hits

df <- df %>%
  mutate(
    HC_label = HC_rate >= threshold & hc_gate,
    WF_label = WF_rate >= threshold,
    CH_label = CH_rate >= threshold,
    SC_label = SC_rate >= threshold,
    MO_label = MO_rate >= threshold,
    n_labels = HC_label + WF_label + CH_label + SC_label + MO_label
  )

# Summary tables

table_primary <- df %>%
  count(primary_type) %>%
  mutate(share = round(n / sum(n) * 100, 1)) %>%
  arrange(desc(n))

```

```
table_function <- df %>%
  count(intervention_function) %>%
  mutate(share = round(n / sum(n) * 100, 1)) %>%
  arrange(desc(n))
```

```
table_logic_function <- df %>%
  count(primary_type, intervention_function) %>%
  tidyr::pivot_wider(
    names_from = intervention_function,
    values_from = n,
    values_fill = 0
  )
```

```
table_secondary <- df %>%
  mutate(secondary_type = ifelse(is.na(secondary_type), "None",
secondary_type)) %>%
  count(primary_type, secondary_type) %>%
  tidyr::pivot_wider(
    names_from = secondary_type,
    values_from = n,
    values_fill = 0
  )
```

```
table_hc <- df %>%
  filter(primary_type == "Human Capital") %>%
  count(HC_subtype) %>%
  mutate(share = round(n / sum(n) * 100, 1)) %>%
  arrange(desc(n))
```

```
table_logic_function_pct <- df %>%
  count(primary_type, intervention_function) %>%
```

```

group_by(primary_type) %>%
mutate(share = round(100 * n / sum(n), 1)) %>%
ungroup() %>%
select(-n) %>%
tidyr::pivot_wider(
  names_from = intervention_function,
  values_from = share,
  values_fill = 0
)

```

```

table_secondary_pct <- df %>%
  mutate(secondary_type = ifelse(is.na(secondary_type), "None",
secondary_type)) %>%
  count(primary_type, secondary_type) %>%
  group_by(primary_type) %>%
  mutate(share = round(100 * n / sum(n), 1)) %>%
  ungroup() %>%
  select(-n) %>%
  tidyr::pivot_wider(
    names_from = secondary_type,
    values_from = share,
    values_fill = 0
  )

```

```

secondary_dist <- df %>%
  filter(!primary_type %in% c("Mixed/Hybrid", "Other/Unclear")) %>%
  mutate(secondary_type = ifelse(is.na(secondary_type), "None",
secondary_type)) %>%
  count(secondary_type) %>%
  mutate(share = round(100 * n / sum(n), 1)) %>%
  tidyr::complete(

```

```

secondary_type = c(
  "Human Capital", "Work-First", "Capability/Health",
  "Social Capital", "Motivational", "None"
),
fill = list(n = 0, share = 0)
) %>%
arrange(desc(n))

table_primary_function_long <- df %>%
  count(primary_type, intervention_function) %>%
  group_by(primary_type) %>%
  mutate(share = round(n / sum(n) * 100, 1)) %>%
  ungroup() %>%
  arrange(primary_type, desc(n))

table_primary_secondary_long <- df %>%
  mutate(secondary_type = ifelse(is.na(secondary_type), "None",
secondary_type)) %>%
  count(primary_type, secondary_type) %>%
  group_by(primary_type) %>%
  mutate(share = round(n / sum(n) * 100, 1)) %>%
  ungroup() %>%
  arrange(primary_type, desc(n))

# Print outputs

cat("\n--- PRIMARY TYPE DISTRIBUTION ---\n")
print(table_primary, n = 50)

cat("\n--- SECONDARY TYPE DISTRIBUTION ---\n")

```

```

print(secondary_dist, n = 50)

cat("\n--- HC SUBTYPES (PRIMARY HC ONLY) ---\n")
print(table_hc, n = 50)

cat("\n--- INTERVENTION FUNCTION DISTRIBUTION ---\n")
print(table_function, n = 50)

cat("\n--- CROSSTABS: PRIMARY TYPE x INTERVENTION FUNCTION ---
\n")
print(table_primary_function_long, n = 100)

cat("\n--- PRIMARY TYPE x INTERVENTION FUNCTION (WIDE) ---\n")
print(table_logic_function, n = 50)

cat("\n--- CROSSTABS: PRIMARY x SECONDARY LOGIC ---\n")
print(table_primary_secondary_long, n = 100)

# Outputs and exports

df_export <- df %>% select(-text)
write_xlsx(df_export, file.path(OUT_DIR, "PO2_classification.xlsx"))
write_csv(df_export, file.path(OUT_DIR, "PO2_classification.csv"))

# Summary tables saved
write_csv(table_primary, file.path(OUT_DIR, "table_primary.csv"))
write_csv(secondary_dist, file.path(OUT_DIR,
"table_secondary_distribution.csv"))
write_csv(table_function, file.path(OUT_DIR, "table_function.csv"))
write_csv(table_hc, file.path(OUT_DIR, "table_hc_subtypes.csv"))

```

```

# Crosstabs saved
write_csv(table_logic_function, file.path(OUT_DIR,
"table_logic_function.csv"))
write_csv(table_secondary, file.path(OUT_DIR, "table_secondary.csv"))
write_csv(table_primary_function_long, file.path(OUT_DIR,
"table_primary_function_long.csv"))
write_csv(table_primary_secondary_long, file.path(OUT_DIR,
"table_primary_secondary_long.csv"))

# Percentage tables
write_csv(table_logic_function_pct, file.path(OUT_DIR,
"table_logic_function_pct.csv"))
write_csv(table_secondary_pct, file.path(OUT_DIR,
"table_secondary_pct.csv"))

cat(
  "\nClassification complete and exported.\n",
  "Outputs saved to: ", OUT_DIR, "\n",
  "Excluded log: ", file.path(OUT_DIR, "excluded_files.csv"), "\n",
  sep = ""
)

# Diagnostics

df_diag <- df %>%
  mutate(row_signal = rowMeans(type_matrix, na.rm = TRUE))

cat("\n--- MEAN ROW SIGNAL + TOP1 BY PRIMARY ---\n")
df_diag %>%

```

```

group_by(primary_type) %>%
summarise(
  n = n(),
  mean_row_signal = mean(row_signal, na.rm = TRUE),
  mean_top1 = mean(top1, na.rm = TRUE),
  .groups = "drop"
) %>%
arrange(desc(n)) %>%
print(n = 50)

diag_logic <- function(type_label) {
  df %>%
    filter(primary_type == type_label) %>%
    summarise(
      n = n(),
      mean_top1 = mean(top1, na.rm = TRUE),
      mean_top2 = mean(top2, na.rm = TRUE),
      mean_dom_finite = mean(dominance_ratio_finite, na.rm = TRUE),
      median_dom_finite = median(dominance_ratio_finite, na.rm = TRUE),
      pct_close_15 = mean(dominance_ratio < 1.5, na.rm = TRUE),
      pct_close_13 = mean(dominance_ratio < 1.3, na.rm = TRUE)
    )
}

cat("\n--- DIAGNOSTICS BY PRIMARY CATEGORY ---\n")
bind_rows(
  diag_logic("Human Capital"),
  diag_logic("Work-First"),
  diag_logic("Capability/Health"),
  diag_logic("Social Capital"),
  diag_logic("Motivational"),

```

```
diag_logic("Mixed/Hybrid"),  
diag_logic("Other/Unclear")  
) %>%  
  arrange(desc(n)) %>%  
  print(n = 50)
```